Decision Analysis

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

On (Measurable) Multiattribute Value Functions: An Expository Argument

James E. Smith, James S. Dyer

http://pubsonline.informs.org/journal/deca

DECISION ANALYSIS
Vol. 18, No. 4, December 2021, pp. 247–256
ISSN 1545-8490 (print), ISSN 1545-8504 (online)

On (Measurable) Multiattribute Value Functions: An Expository Argument

James E. Smith,[a] James S. Dyer[b]

[a] Tuck School of Business, Dartmouth College, Hanover, New Hampshire 03755; [b] McCombs School of Business, The University of Texas at Austin, Austin, Texas 78712

Contact: jim.smith@dartmouth.edu, https://orcid.org/0000-0002-9429-7567 (JES); jim.dyer@mccombs.utexas.edu (JSD)

Received: October 13, 2020
Revised: April 28, 2021
Accepted: May 8, 2021
Published Online in Articles in Advance: August 6, 2021

https://doi.org/10.1287/deca.2021.0435

Abstract. In this note, we provide an easy-to-understand introduction to strength-of-preference measures in the context of deterministic multiattribute value assessments, focusing on what they are and why they matter. Though these issues are well understood by some decision analysts, we believe that many do not understand or appreciate the role of strength-of-preference assumptions when assessing or interpreting multiattribute value functions. The note is structured around an argument between the two authors that took place when reviewing applications of multiattribute value functions.

Keywords: multiattribute value theory

1. Introduction

For the 65th anniversary of *Management Science*, we wrote a paper reviewing the impact of decision analysis research in *Management Science* on practice (Dyer and Smith 2020). As part of this, we reviewed several applications of multiattribute value theory, where the standard model is an additive value function. As an example, we highlighted Parnell et al. (1998) who developed an additive value function with 134 attributes to evaluate 43 systems concepts to support the goal of the US Air Force to achieve space and air dominance by 2025. In this study, the analysts had the project teams agree on the shapes of the attribute-specific value functions (linear, concave, convex, or s-curve). These value functions were then assessed and assigned weights and combined in an additive value function.

In a draft of the paper, **Dyer** called the value function developed in Parnell et al. (1998) a *measurable* additive value function. However, **Smith** objected to this characterization because Parnell et al. made no mention of measurability of the value function; Parnell et al. said only that they used "standard decision analysis techniques" (p. 1344).

- **Smith:** It is an additive value function. Isn't mutual preferential independence (as discussed in Keeney and Raiffa 1976) enough to ensure additivity?
- **Dyer:** Yes, but, any time the attribute value functions are defined independently without using a method involving explicit tradeoffs with another attribute, the implicit assumption is that it is a measurable value function built on strength-of-preference assessments (as in Dyer and Sarin 1979).
- **Smith:** But people do this all the time, without mentioning measurability or strength-of-preference relations!
- **Dyer** (patiently): Yes, they do.
- **Smith** (skeptically): But do we really need to make assumptions about strength-of-preference relations to justify this practice? Couldn't we get to this form using willingness-to-pay arguments instead?
- **Dyer**: Well … that's a longer story.

This back and forth convinced us that the concept of measurability and its role in multiattribute value modeling may not be well understood in the decision analysis community. The theory of measurable value functions was an active research area in decision analysis in the

1970s and 1980s and related issues are carefully discussed in some textbooks (e.g., French 1986, von Winterfeldt and Edwards 1986, Kirkwood 1997). However, measurability is not mentioned at all in other popular decision analysis textbooks that discuss multiattribute value functions (e.g., Keeney and Raiffa 1976, Clemen and Reilly 2013). Of course, there are many examples of papers describing applications of multiattribute value models that address measurability issues; examples from this journal include Ewing et al. (2006), Bana e Costa et al. (2008), and Dees et al. (2013). But there are other papers that justify the use of an additive value function by citing mutual preferential independence assumptions without mentioning the related strength-of-preference assumptions; examples from this journal include Merrick et al. (2005), Simon and Melese (2011), and De Icaza et al. (2019). There are also many other papers—like Parnell et al. (1998)—that only say they use "standard decision analysis techniques" to assess an additive value function. These examples and others made us think that **Smith** is not alone in not fully appreciating the role that measurability assumptions play in multiattribute value modeling.

The purpose of this note is to provide an easy-to-understand introduction to measurability in the context of (deterministic) multiattribute value assessments, focusing on what it is and why it matters. In particular, we describe how "standard decision analysis techniques" for assessing additive value functions assume mutual preferential independence along with two additional conditions, difference independence and difference consistency, which relate to measurability. We structure the note as an elaboration of the argument between **Dyer** and **Smith** outlined above.

2. Measurable Value Functions

To illustrate measurable value functions, we will use an example from von Winterfeldt and Edwards (1986; §7.3) that considers the location preferences of a new medical doctor considering offers for positions for an assistant professorship. The example focuses on her location preferences (reflecting proximity to family and friends, climate, and culture, among other factors) and sets aside other aspects of the offers (salary, prestige, etc.). In the example, the analyst asks her to rank five locations and then assign numeric values to reflect her preferences. Her most preferred location (San Francisco) is assigned a value of 100 and her least preferred location (Ann Arbor) is assigned a value of 0. The other locations (Boston, Los Angeles, and Chicago) have values that lie between these two extremes; see Table 1. These value judgments reflect her *preferences* (with higher values corresponding to more preferred locations) and her *strength of preferences*. For example, San Francisco is strongly preferred to the other locations, whereas Chicago and Ann Arbor are closer in her view.

The value function of Table 1 is an example of *measurable value function* v, which represents both a preference order $\succcurlyeq$ on alternatives and a strength-of-preference order $\succcurlyeq^*$ on pairs of alternatives,[1] such that:

$$x \succcurlyeq y \;\; \textit{if and only if} \;\; v(x) \geq v(y) \;\; \text{and} \tag{1a}$$

$$(x', y') \succcurlyeq^* (x'', y'') \;\; \textit{if and only if} \;\; v(x') - v(y') \geq v(x'') - v(y'') \tag{1b}$$

Here the first relation (1a) captures preferences in the obvious way: given a choice between alternatives x and y, the decision maker (DM) prefers the alternative with the larger value of v. The second relation (1b) is perhaps less obvious and can be interpreted as meaning that the DM thinks that it is a "bigger improvement" to go to x' from y' than to go to x'' from y'' or that the DM would prefer the "exchange" of x' for y' to the exchange of x'' for y''. For instance, in the example of Table 1, the DM finds the improvement in going from Boston to San Francisco to be "more of an improvement" than that of going from Ann Arbor to Chicago. This captures the intuitive idea that San Francisco is much preferred to Boston (and the other cities), but Ann Arbor and Chicago are "close."

Note that any increasing transformation of the value function v preserves the order of values and thus also represents the preference relation $\succcurlyeq$ (i.e., is strategically equivalent), whereas only positive linear

Table 1. A Measurable Value Function for Location Preferences

Location	Value
San Francisco	100
Boston	60
Los Angeles	40
Chicago	15
Ann Arbor	0

transformations of v would also represent both the preference relation $\succcurlyeq$ and the strength-of-preference relation $\succcurlyeq^*$. The theory underlying these measurable value functions is developed in detail in Krantz et al. (1971; §4) and French (1986; §9), and is summarized in a very accessible form in von Winterfeldt and Edwards (1986; §9.2).

A measurable value function can also be defined on a continuous range of outcomes. As suggested by Kirkwood (1997; p. 66), the procedure for assessing such a value function may be based on asking the DM to identify the *midvalue* for a range of outcomes. "The midvalue of a range is defined to be the score such that the difference in value between the lowest outcome score in the range and the midvalue is the same as the difference in value between the midvalue and the highest score." If the endpoints of the range are scaled to have values 0 and 1, the midvalue would have a value of 0.5. This process can continue by assessing the midvalues of the segments above and below the first midvalue, and so on until enough points are assessed to allow the approximation of the value function.

Clearly, the strength-of-preference notion has some intuitive appeal: we understand what the doctor means when she says San Francisco is "strongly preferred" to the other locations. However, despite this intuitive appeal, the concept of a measurable value function and strength-of-preference measures have been viewed with skepticism by some economists, decision theorists, and decision analysts for many years (not just **Smith**). Many of the concerns focus on what exactly a strength-of-preference means: it is not clear what a bigger improvement is or what preferences for exchanges are. For example, Machina (1981, p. 169) notes that if he were asked to compare the strength of preference for one improvement over another, he would "respond to this question by asking what it meant." Underlying these concerns is the fact that strength-of-preference judgments are "nonoperational" because they do not correspond to any real or hypothetical choice behavior. Related concerns focus on why it is necessary: If the goal is to help DMs make choices, why do we need to consider strength-of-preference measures? See Farquhar and Keller (1989), von Winterfeldt and Edwards (1986; §7.1), and French (1986; §9) for more discussion of these criticisms and responses.

3. Additive Multiattribute Value Functions: Theory

As discussed in the introduction, the most common approach for evaluating multiattribute alternatives is to use an additive representation. Let $\boldsymbol{x} = (x_1, \ldots, x_n)$ denote a multiattribute alternative with n attributes $X_1, \ldots, X_n$. An additive value function has the form

$$v(x_1, \ldots, x_n) = \sum_{i=1}^{n} v_i(x_i) \tag{2}$$

where $v_i(x_i)$ is an attribute-specific value function for X_i. If we let x_i^* and x_i^0 denote the best and worst levels for attribute X_i, we can normalize the additive value function (2) by taking

$$v(x_1, \ldots, x_n) = \sum_{i=1}^{n} \lambda_i \, v_i(x_i) \tag{3}$$

where $v_i(x_i^0) = 0$, $v_i(x_i^*) = 1$, and the weights λ_i are scaled to be between 0 and 1 with $\sum_{i=1}^{n} \lambda_i = 1$. This implies that $v(x_1^0, \ldots, x_n^0) = 0$ and $v(x_1^*, \ldots, x_n^*) = 1$.

If $n \geq 3$, an additive value function exists if and only if the attributes are *mutually preferentially independent* (see, e.g., Debreu 1960, Keeney and Raiffa 1976, p. 111).[2] To define mutual preferential independence, let $I \subseteq \{1, \ldots n\}$ denote a subset of the attribute indices and define X_I as the subset of the attributes with indices in I; let $\bar{X}_I$ be the complementary subset of the attributes. Then, X_I is *preferentially independent* of $\bar{X}_I$ if

$$(y_I, \bar{x}_I) \succcurlyeq (x_I, \bar{x}_I) \text{ for any } x_I, y_I \in X_I \text{ and } \bar{x}_I \in \bar{X}_I \text{ implies}$$
$$(y_I, \bar{y}_I) \succcurlyeq (x_I, \bar{y}_I) \text{ for all } \bar{y}_I \in \bar{X}_I.$$

Intuitively, this means that preferences for attributes in X_I (i.e., tradeoffs among these attributes) do not depend on the levels of the complementary attributes $\bar{X}_I$. The attributes $X_1, \ldots, X_n$ are mutually preferentially independent if every subset of attributes X_I is preferentially independent of its complementary subset of attributes $\bar{X}_I$. The resulting additive form is unique in that any two additive value functions that represent the same ordinal preferences ($\succcurlyeq$) must be related by a positive linear transformation; thus the normalized additive form (3) is unique.

Note that these requirements for an additive value function concern only the DM's preferences (captured by the $\succcurlyeq$ relation) and do not reference strength-of-preferences (captured by the $\succcurlyeq^*$ relation). Thus, as noted by **Smith** in the argument in the introduction, the use of the additive value function rests on mutual preferential independence, without any appeal to measurability or strength-of-preference assumptions.

4. Additive Multiattribute Value Functions: Assessment

Yes, says **Dyer**, but let us look carefully at this conclusion: the result says that, given mutual preferential independence, a unique additive value function *exists*. But we need to perform careful assessments to determine the attribute-specific value functions $v_i(x_i)$ as well as the weights λ_i. These assessment procedures are described in detail in Keeney and Raiffa (1976, §3.7) and Kirkwood (1997; §9.2). Given the attributes, the process consists of three steps:

1. Confirm (or assume) the attributes satisfy mutual preferential independence;
2. Determine the attribute-specific value functions $v_i(x_i)$; and
3. Determine the weights λ_i.

We will focus on step (2), determining the attribute-specific value functions $v_i(x_i)$ using the *midvalue splitting technique*, which seems analogous to the univariate midvalue assessment technique discussed in §2, but, as we will see, is much more involved.

Given a normalized additive value function of the form of (3), Kirkwood (1997; p. 233) defines the midvalue of an interval as follows:

> *The* midvalue of an interval $[x'_i, x''_i]$ *is the level x_i^m such that starting from a specified level of another attribute, the DM would give up the same amount of that other attribute to improve from x'_i to x_i^m as to improve from x_i^m to x''_i.*

The attribute-specific value of this midvalue $v_i(x_i^m)$ is then equal to $0.5\, v_i(x'_i) + 0.5\, v_i(x''_i)$. In the midvalue splitting technique, the analyst begins by assessing the midvalue $x_i^{0.50}$ of the interval from best to worst levels of the attribute $[x_i^0, x_i^*]$. Given the normalization of the value function, we have $v_i(x_i^{0.50}) = 0.5\, v_i(x_i^0) + 0.5\, v_i(x_i^*) = 0.5$. Next, the analyst assesses the midvalue $x_i^{0.75}$ of the interval $[x_i^{0.5}, x_i^*]$ and the midvalue $x_i^{0.25}$ of the interval $[x_i^0, x_i^{0.5}]$, having values $v(x_i^{0.75}) = 0.75$ and $v(x_i^{0.25}) = 0.25$. At this point, the analyst has five points of the attribute-specific value function $v_i(x_i)$ and can subdivide the intervals further to achieve greater refinement, if so desired.

Given mutual preferential independence, an additive value function exists and there is no assumption that the attribute-specific value functions represent a strength-of-preference relationship. Nevertheless, the midvalue assessment question

> *What value x_i^m yields the same improvement from x'_i to x_i^m as the improvement from x_i^m to x''_i?*

or, phrased differently,

> *What value x_i^m yields half of the improvement of going from x'_i to x''_i?*

is very similar to a strength-of-preference assessment question. It is easy and natural to ignore the change in values of the "other attribute" when answering these questions. Indeed, Kirkwood (1997), after defining the midvalue as above, makes no further mention of the other attribute when discussing the midvalue assessment technique. Without a strength-of-preference interpretation and additional independence assumptions about strength-of-preferences, these midvalue assessment questions can be very difficult to answer correctly and may be easily misconstrued.

To illustrate the difficulty of the midvalue assessment technique, consider an example of a value function for quality of life years that includes consideration of health state and consumption. For simplicity, we focus on a value function where q denotes a quality of life measure (scaled between 0 for a health state comparable with death and 1 for perfect health), l denotes the number of years lived in this health state, and c the annual consumption over this lifespan.[3] Suppose the qs for the problem under consideration range from $q^0 = 0.2$ to $q^* = 1.0$ and the ls range from $l^0 = 0$ years to $l^* = 50$ years. We will assume that, unbeknownst to the analyst, the DM has a measurable value function of the form

$$v(q, l, c) = q \times l^{\eta} \times u(c), \tag{4}$$

where $\eta = 0.50$. We will not make any specific assumptions about the form of $u(c)$ other than that it is increasing and positive over the range of consumption levels considered. Notice that this DM's preferences satisfy mutual preferential independence (preferences for any two attributes do not depend on a common value of the third attribute). Taking logarithms, the

DM's ordinal preference relation ($\succcurlyeq$) can be represented in an additive form as

$$\ln(q) + \eta\ln(l) + \ln(u(c)). \quad (5)$$

If the analyst asks the DM midvalue questions—What is the midvalue of q in the range $[0.2, 1]$?—it would be natural for the DM to appeal to the strength-of-preference interpretation when answering these questions and ignore changes in the other attribute. Based on (4), for q this would lead to the assessments that the midvalue of $[0.2, 1]$ is $q^{0.50} = 0.6$, the midvalue of $[0.6, 1]$ is $q^{0.75} = 0.8$ and so on, leading to an attribute specific value function $v_q(q) \propto q$. Following a similar process for l and c, the analyst would find attribute-specific value functions $v_l(l) \propto l^\eta$ and $v_c(c) \propto u(c)$ and incorrectly arrive at a value function of the form

$$\lambda_q q + \lambda_l l^\eta + \lambda_c u(c), \quad (6)$$

rather than a correct form such as (4) or (5). Note that (5) is an order-preserving transformation of (4) and both would provide the same rankings of alternatives, whereas the additive model (6) is not an order-preserving transformation of (4) and may produce incorrect rankings.

Properly assessing the value function requires paying careful attention to the other attribute in the definition of the midvalue. Suppose we seek to assess $v_q(q)$ and take a base value of the other attribute l to be $l^* = 50$ years; we will assume consumption is held constant at c. To find the midvalue of $q^{0.50}$ in the range $[q^0, q^*]$, the definition of the midvalue—*read carefully*—requires the DM to determine a $q^{0.50}$ and the amount of that other attribute l given up δ_l that satisfy the following two indifference conditions:

$$(q^0, l^*, c) \sim (q^{0.50}, l^* - \delta_l, c) \quad (7a)$$

$$(q^{0.50}, l^*, c) \sim (q^*, l^* - \delta_l, c) \quad (7b)$$

The required assessment is illustrated in Figure 1: given the point (q^0, l^*) (labeled point 1 in the figure), we seek a point $(q^{0.50}, l^* - \delta_l)$ (labeled point 2) on the same indifference curve, such that the DM is also indifferent between $(q^{0.50}, l^*)$ (labeled point 3) and $(q^*, l^* - \delta_l)$ (labeled point 4), having the same quality $q^{0.50}$ and lifetime value $l^* - \delta_l$ (respectively) as point 2. Given a normalized additive value function (3), these conditions are equivalent to

Figure 1. (Color online) Indifference Curves and Points Involved in the Midvalue Assessment

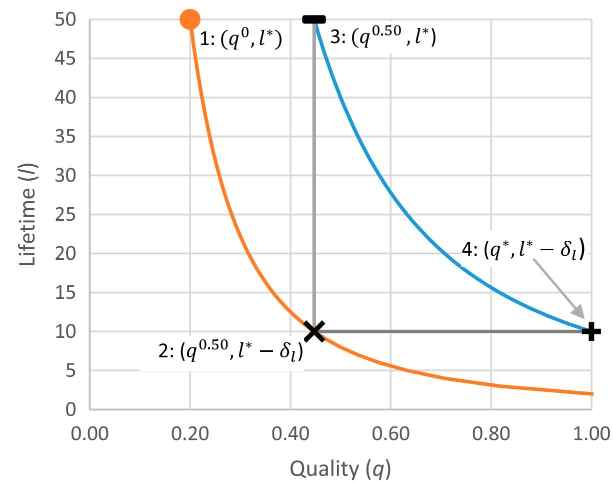

$$\begin{aligned}\lambda_q v_q(q^0) &+ \lambda_l v_l(l^*) + \lambda_c v_c(c) \\ &= \lambda_q v_q(q^{0.50}) + \lambda_l v_l(l^* - \delta_l) + \lambda_c v_c(c)\end{aligned} \quad (8a)$$

$$\begin{aligned}\lambda_q v_q(q^{0.50}) &+ \lambda_l v_l(l^*) + \lambda_c v_c(c) \\ &= \lambda_q v_q(q^*) + \lambda_l v_l(l^* - \delta_l) + \lambda_c v_c(c)\ .\end{aligned} \quad (8b)$$

Subtracting (8b) from (8a), we find the midvalue satisfies

$$v_q(q^{0.50}) = 0.5\ v_q(q^0) + 0.5 v_q(q^*) = 0.50.$$

This result is apparently independent of the assumed base for l (here taken to be l^*) and the amount of that other attribute given up δ_l, though these values are critical to the definition of the midvalue $q^{0.50}$.

The assessment of the indifference points $q^{0.50}$ and δ_l satisfying conditions (7a) and (7b)—requiring simultaneous consideration of two indifference curves—is likely to be very difficult. With a true value function of the form (4), we can solve two nonlinear equations (corresponding to these indifference conditions) in two unknowns ($q^{0.50}$ and δ_l) to determine the midvalue:

$$v(q^0, l^*, c) = v(q^{0.50}, l^* - \delta_l, c) \quad (9a)$$

$$v(q^{0.50}, l^*, c) = v(q^*, l^* - \delta_l, c) \quad (9b)$$

Given our specific numerical assumptions, we can solve these two equations (e.g., by using Solver in Excel) to find $q^{0.50} = 0.4472$ and $\delta_l = 40.0$. Recall that the natural, but incorrect, strength-of-preference interpretation gave the midvalue $q^{0.50} = 0.60$.[4]

If we continue to assess the attribute-specific value function v_q using this midvalue splitting technique, use the same procedure to assess v_l and v_c, and then properly assess the scaling weights λ_q, λ_l, and λ_c (e.g., following the approach recommended in Keeney and Raiffa (1976) or Kirkwood (1997)), we arrive at a normalized additive value function

$$\frac{1}{2+\eta}\frac{\ln(q)-\ln(q^0)}{\ln(q^*)-\ln(q^0)}+\frac{\eta}{2+\eta}\frac{\ln(l)-\ln(l^0)}{\ln(l^*)-\ln(l^0)} + \frac{1}{2+\eta}\frac{\ln(u(c))-\ln(u(c^0))}{\ln(u(c^*))-\ln(u(c^0))}. \tag{10}$$

This value function is strategically equivalent to (4) and (5) and thus correctly represents the DM's ordinal preferences ($\succcurlyeq$) but it does not correctly represent the DM's strength of preferences ($\succcurlyeq^*$).

Our ability to interpret the shapes of attribute-specific value functions also rests on the correct representation of the strength of preference relation. For example, Parnell et al. (1998, p. 1345) "had the team agree on the shape of the curve (linear, concave, convex or s-curve)" before assessing points on the curve. But such conclusions about the shape of the curve are not meaningful as these shapes may not be preserved under increasing transformations of the measurable value function even if the decision maker's preferences are mutually preferentially independent. For example, holding other attributes constant, the measurable value function (4) is linear in q, whereas the attribute-specific value function in the order-preserving additive form (10) is proportional to $\ln(q)$. (The "meaningfulness" of preference statements is discussed in detail in Krantz et al. (1971) and French (1986; §9.2).)

Thus, argues **Dyer**, there is nothing wrong with the theory of additive value functions or applications with a careful assessment process. But in practice, almost no one does this because it would be very difficult for a DM to make the required tradeoffs. Indeed the only application of the midvalue splitting technique that we are aware of is the work by Roche, which is described in Keeney and Raiffa (1976, §7.2).[5] Instead, most practitioners using "standard decision analysis techniques" assume—implicitly or otherwise—that the value function is a measurable additive value function, satisfying the assumptions described below. These assumptions allow one to assess attribute-specific value functions without considering tradeoffs with other attributes. If these assumptions are not satisfied and one uses standard decision analysis techniques, one may arrive at an incorrect representation of the value function, such as (6).

5. Measurable Multiattribute Value Functions

The theory of measurable multiattribute value functions presumes the existence of a strength-of-preference order $\succcurlyeq^*$ as well as a preference order $\succcurlyeq$ on alternatives and makes assumptions about both. An additive value function of the form of Equation (2) or (3) is measurable if it represents both a preference order $\succcurlyeq$ on alternatives and a strength-of-preference order $\succcurlyeq^*$, that is, satisfies conditions (1a) and (1b). What is required to ensure this result? Following Dyer and Sarin (1979), we can identify three conditions on the preference relation $\succcurlyeq$ and strength-of-preference relation $\succcurlyeq^*$ that are sufficient for this purpose:

1. the attributes are mutually preferentially independent, as discussed earlier;
2. the attributes are *difference consistent*; and
3. one of the attributes is *difference independent* of the others.

Difference consistency means that

$$\textit{If } x \succcurlyeq y \succcurlyeq z, \textit{ then } (x,z) \succcurlyeq^* (y,z),$$

which means that if the DM prefers x to y, then for any less preferred outcome z, the DM views going from z to x as a bigger improvement (or preferred exchange) than going from z to y (x, y, and z here are multiattribute alternatives). As the name suggests, this is a basic consistency requirement that relates the preference relation $\succcurlyeq$ and the strength-of-preference relation $\succcurlyeq^*$.

Attribute X_i is difference independent of the others if

$$(x_i,\bar{x}_i) \succcurlyeq (y_i,\bar{x}_i) \textit{ for some } x_i, y_i \in X_i \text{ and } \bar{x}_i \in \bar{X}_i \textit{ implies}$$
$$(x_i,\bar{x}_i)(y_i,\bar{x}_i) \sim^* (x_i,\bar{y}_i)(y_i,\bar{y}_i) \textit{ for any } \bar{y}_i \in \bar{X}_i.$$

Intuitively, this means the improvement in going from y_i to x_i on one attribute (all other attributes held constant) is the same regardless of the levels of the other attributes. This condition relates directly to the strength-of-preference relationship and asks the DM to consider preferences for improvements; the condition ensures that the value associated with changes in one attribute does not depend on the levels of the other attributes.

With these assumptions, one can assess measurable attribute-specific value functions using the midvalue assessment questions of the form considered earlier

> *What value x_i^m yields half of the improvement of going from x_i' to x_i''?*

or, alternatively,

> *Rate Boston on a scale from 0 (Ann Arbor) to 100 (San Francisco)*

without explicitly considering tradeoffs with other attributes. Similarly, one can talk about the shape of these attribute-specific value functions—for example, are they linear, concave, convex, or s-shaped? It is the measurability of the value function—ensured by the difference consistency and independence assumptions paired with mutual preferential independence—that allows these simplifications in assessing and interpreting the standard additive value function. These assumptions are also implicit in popular computer programs that support the assessment of additive value functions, including Logical Decisions (Logical Decisions 2021) and Hiview3 (Catalyze Ltd 2021). These are things that people "do all the time," as **Dyer** and **Smith** noted in the introduction.

In the quality of life example of Equation (4), the mutual preferential independence condition held but the difference independence condition did not—for example, improvements in health state are valued more given a longer life—and this caused "standard decision analysis techniques" to lead to the incorrect representation (6). If preferential independence holds, but difference independence does not, one can sometimes appeal to a *weak difference independence* condition, which, when combined with mutual preferential independence, justifies the existence of either an additive or a multiplicative measurable value function such as (4). This weak indifference independence condition requires the rank order of preference differences (but not their magnitudes) to be independent of the levels of the other attributes; see Dyer and Sarin (1979), theorem 3. Alternatively, one might redefine attributes so difference independence holds with the new attributes, analogously to Keeney (1981)'s methods for ensuring additive independence with multiattribute utility functions.

6. Willingness to Pay and Measurability

Given the skepticism about strength-of-preference measurements discussed in §1 is there some other more concrete way to justify this standard practice for assessing additive value functions, perhaps using willingness-to-pay assumptions as **Smith** suggested?

To illustrate the willingness-to-pay interpretation, let us expand the example of the medical doctor choosing among assistant professor offers discussed in §2 to consider salary (s) and prestige (p) attributes, as well as location (l). In the earlier example, preferences for different locations were measured using values ranging from 0 to 100 to represent strength-of-preference for the five cities in Table 1. Salary might be measured in terms of cost-of-living-adjusted dollars per year and prestige might be scored on a discrete scale; for example, using letter grades from "A" (the maximum possible prestige) to "F." We also assume that the DM strictly prefers larger salaries to smaller salaries, all else equal.

In the willingness-to-pay approach (see Keeney and Raiffa 1976, §3.8), one chooses (or introduces) a "money" attribute and compares alternatives to base-case levels of the other attributes. In the example, we will take salary s to be the money attribute and take the base case for the other attributes to be (l^0, p^0). To find the willingness-to-pay for (s, l, p), we seek a Δ such that

$$(s, l, p) \sim (s + \Delta,\ l^0, p^0)\ .$$

The interpretation is that the DM is indifferent to accepting an increment Δ in salary (from s) and having (l^0, p^0) instead of (l, p). Using the assumption that the DM strictly prefers larger salaries, we could then rank alternatives in terms of their willingness-to-pay adjusted salaries, $s + \Delta$.[6] We can price out changes in different attributes one attribute at a time as in the "even swaps" method proposed by Hammond et al. (1999). This pricing out approach does not require any specific assumptions about the form of the preference model, but, without more structure on the DM's preferences, the assessment of such willingness-to-pay values is likely to be difficult if there are many alternatives to consider.

To simplify these willingness-to-pay assessments, two assumptions are standard. First, it is standard to assume that the attributes are mutually preferentially independent. This assumption allows the DM to consider, for example, tradeoffs between salary s and location l without simultaneously considering prestige p. Thus, we can decompose the willingness-to-pay Δ into attribute-specific willingness-to-pay amounts. For example, let Δ_l be defined by the following indifference assessment:

$$(s, l, p) \sim (s + \Delta_l,\ l^0, p).$$

For example, suppose the DM considers an alternative with a cost-of-living adjusted salary of \$100k located in Los Angeles with a prestige rating of p and would be just indifferent to a cost-of-living adjusted salary of \$120k in Ann Arbor (l^0) with the identical prestige rating; then Δ_l is \$20k. A similar tradeoff between salary and prestige would complete the process:

$$(s + \Delta_l, l^0, p) \sim (s + \Delta_l + \Delta_p,\ l^0, p^0)$$

Here the total willingness-to-pay is decomposed as $\Delta = \Delta_l + \Delta_p$ where Δ_l and Δ_p are the amounts the DM is willing to pay to go from location l to l^0 and then from prestige p to p^0. Mutual preferential independence allows Δ_l to depend on s but ensures that Δ_l does not depend on p. Similarly, Δ_p may depend on $s + \Delta_l$ but Δ_p does not depend on l.

The second standard simplifying assumption in willingness-to-pay assessments is to assume that the attribute-specific willingness-to-pays (here Δ_l and Δ_p) do not depend on the level of the money attribute (here salary s). This, coupled with the mutual preferential independence assumption, allows the analyst to assess the attribute-specific willingness-to-pay functions $\Delta_l(l)$ and $\Delta_p(p)$ independently and sum these terms to form a value function. Thus, these willingness-to-pay arguments lead to an additive value function

$$v(\mathrm{s}, l, p) = s + \Delta_l(l) + \Delta_p(p), \tag{9}$$

where the attribute-specific willingness-to-pay functions $\Delta_l(l)$ and $\Delta_p(p)$ are measured in salary and play the role of attribute-specific value functions. Thus, **Smith** argues, this willingness-to-pay-based value function has a concrete interpretation in terms of the DM's preferences ($\succcurlyeq$) and it is meaningful to talk about the attribute-specific functions being convex, concave, or s-shaped, for example. This interpretation does not require any assumptions about a strength-of-preference relation $\succcurlyeq^*$.

Yes, says **Dyer**, but note that (9) is a measurable additive value function: $v(\mathrm{s}, l, p)$ represents the DM's preferences $\succcurlyeq$ and the location and prestige adjustments to salaries define a strength-of-preference measure defining the relation $\succcurlyeq^*$. In making these willingness-to-pay assessments, we have (i) assumed mutual preferential independence (the first simplifying assumption) and (ii) difference independence (the second simplifying assumption). Difference consistency is satisfied by construction. Thus, this willingness-to-pay assessment process implicitly makes the measurability assumptions required to justify the additive measurable value function. However, this willingness-to-pay approach is restrictive in that the second assumption requires the value function to be linear in the money attribute; this is not necessary in general. For example, one could imagine the attribute-specific value function $v_s(s)$ being concave in salary; this would imply that the doctor's willingness-to-pay for living in San Francisco would be larger with higher salaries. Moreover, the assumption that there is a monetary attribute may be unrealistic in some settings; for example, in some public sector applications, the domain expert whose preferences are relevant for some attributes may be quite removed from the funding organization that cares about costs.

Smith: Ah, it is mutual preferential independence and difference independence that allows us to use these willingness-to-pay methods and leads to a *measurable* additive value function. So, with a justification that relies on either strength-of-preference interpretations or willingness-to-pay arguments, we arrive at the same place.

Dyer: Exactly.

7. Conclusions: Implications for Practice

Smith: I like the concrete interpretation of strength-of-preferences associated with the willingness-to-pay approach, but acknowledge the restrictiveness of the simplifying assumptions. Before this exchange, I had not appreciated the role of difference independence in standard decision analysis techniques for assessing additive value functions. This difference independence condition seems important to confirm in applications, along with mutual preferential independence.

Dyer: I agree. One should check:

• Do tradeoffs between two attributes depend on the levels of the other attributes? (mutual preferential independence); and

• Does the improvement associated with changes in one attribute depend on the levels of the other attributes? (difference independence (Dyer and Sarin 1979; §4 has a longer discussion on this point)).

Like many other practitioners and researchers (and despite the objections of other practitioners and researchers), I have found that people are quite willing to think about comparing improvements in attribute values, even when it is not a monetary tradeoff; this reflects an intuitive sense that value is measurable and meaningful for them.

Smith: Yes, people do it all the time. Perhaps the takeaway message is to encourage people to think carefully about what "value" is in the context of multiattribute value modeling, and to take care in choosing assessment techniques that are appropriate for a given application.

Acknowledgments
The authors would like to thank three referees and the associate editor for helpful feedback on this manuscript.

Endnotes

[1] We let $\succ$ $(\succ^*)$ and $\sim$ $(\sim^*)$ denote the strict preference and indifference relations corresponding to $\succcurlyeq (\succcurlyeq^*)$.

[2] When there are only two attributes, the corresponding tradeoffs condition is required for additivity. See Keeney and Raiffa (1976) or Kirkwood (1997) for discussions.

[3] For more detailed models considering health states and consumption levels varying over time, see for example, Smith and Keeney (2005) or Lichtendahl and Bodily (2012).

[4] Keeney and Raiffa (1976) and Kirkwood (1997) observe that for a value function satisfying mutual preferential independence, the resulting midvalue point (here $q^{0.50}$) does not depend on the base value for the other attribute (here assumed to be l^*), but the resulting change in the other attribute (δ_l) may be different with a different base value. For instance, if we take the base value $l = 45$ rather than $l^* = 50$, solving (8a) and (8b), we find $q^{0.50} = 0.447$ (as before) and $\delta_l = 36.0$.

[5] Even in Roche's application of the midvalue assessment technique, it is not clear that subjects explicitly considered tradeoffs with other attributes. Keeney and Raiffa's summary of Roche's work says that he used the midvalue method and concluded with a discussion of "the general shape of the v-component value functions" (Keeney and Raiffa 1976, p. 371). It would hardly be surprising if Roche and his subjects also appealed to the intuitive measurable value function interpretation of the midvalue assessment questions, as described above.

[6] To see this, suppose $(s', l', p') \sim s' + \Delta'$, and $(s'', l'', p'') \sim s'' + \Delta''$. Then the former is preferred to the latter if and only if $s' + \Delta' \geq s'' + \Delta''$.

References

Bana e Costa CA, Lourenço JC, Chagas MP, Bana e Costa JC (2008) Development of reusable bid evaluation models for the Portuguese Electric Transmission Company. *Decision Anal.* 5(1):22–42.

Catalyze, Ltd. (2021) *Hiview3*: Software, Version 3.2 (Winchester, England).

Clemen RT, Reilly T (2013) *Making Hard Decisions* (South-Western, Mason, OH).

De Icaza RR, Parnell GS, Pohl EA (2019) Gulf Coast port selection using multiple-objective decision analysis. *Decision Anal.* 16(2): 87–104.

Debreu G (1960) Topological methods in cardinal utility theory. Arrow KJ, Karlin S, Suppes P, eds. *Mathematical Methods in the Social Sciences, 1959: 1st Stanford Symposium: Selected Papers* (Stanford University Press, Palo Alto, CA).

Dees RA, Nestler ST, Kewley R (2013) WholeSoldier performance appraisal to support mentoring and personnel decisions. *Decision Anal.* 10(1):82–97.

Dyer JD, Sarin RK (1979) Measurable multiattribute value functions. *Oper. Res.* 27(4):810–822.

Dyer JD, Smith JE (2020) Innovations in the science and practice of decision analysis: The role of *Management Science*. *Management Sci.* ePub ahead of print August 4, https://doi.org/10.1287/mnsc.2020.3652.

Ewing PL, Tarantino W, Parnell GS (2006) Use of decision analysis in the army base realignment and closure (BRAC) 2005 military value analysis. *Decision Anal.* 3(1):33–49.

Farquhar PH, Keller LR (1989) Preference intensity measurement. *Ann. Oper. Res.* 19:205–217.

French S (1986) *Decision Theory: An Introduction to the Mathematics of Rationality* (Halsted Press, Sydney).

Hammond JS, Keeney RL, Raiffa H (1999) *Smart Choices: A Practical Guide to Making Better Decisions* (Harvard Business School Press, Boston).

Keeney RL (1981) Analysis of preference dependencies among objectives. *Oper. Res.* 29(6):1105–1120.

Keeney RL, Raiffa H (1976) *Decisions with Multiple Objectives: Preferences and Value Tradeoffs* (Wiley, New York).

Kirkwood CW (1997) *Strategic Decision Making* (Duxbury Press, Belmont, CA).

Krantz DH, Luce RD, Suppes P, Tversky A (1971) Foundations of Measurement, vol. I: Additive and Polynomial Representations (Academic Press, New York).

Lichtendahl KC, Bodily SE (2012) Multiplicative utilities for health and consumption. *Decision Anal.* 9(4):314–328.

Logical Decisions (2021) *Logical Decisions*, Version 7.2 Software, Fairfax, VA.

Machina MJ (1981) "Rational" decision making vs. "rational" decision modelling? *J. Math. Psych.* 24(2):163–175.

Merrick JR, Parnell GS, Barnett J, Garcia M (2005) A multiple-objective decision analysis of stakeholder values to identify watershed improvement needs. *Decision Anal.* 2(1):44–57.

Parnell GS, Conley HW, Jackson JA, Lehmkuhl LJ, Andrew JM (1998) Foundations 2025: A value model for evaluating future air and space forces. *Management Sci.* 44(10):1336–1350.

Simon J, Melese F (2011) A multiattribute sealed-bid procurement auction with multiple budgets for government vendor selection. *Decision Anal*. 8(3):170–179.

Smith JE, Keeney RL (2005) Your money or your life: A prescriptive model for health, safety, and consumption decisions. *Management Sci*. 51(9):1309–1325.

von Winterfeldt D, Edwards W (1986) *Decision Analysis and Behavioral Research* (Cambridge University Press, Cambridge).

James E. Smith is the Jack Byrne Distinguished Professor in Decision Science at the Tuck School of Business at Dartmouth College. His main research interest is in decision analysis, particularly in decision making under uncertainty over time (e.g., dynamic programming). Smith received INFORMS' Frank P. Ramsey Medal for distinguished contributions to decision analysis in 2008. His hobbies include arguing with Jim Dyer.

James S. Dyer is the Fondren Foundation Centennial Chair in Business at the McCombs School of Business at the University of Texas at Austin. His research and teaching interests focus on decision analysis including risk management, multiple criteria decision making, and capital budgeting. Dyer received INFORMS' Frank P. Ramsey Medal for distinguished contributions to decision analysis in 2002. His hobbies include arguing with Jim Smith.