

MARC LE MENESTREL and LUK VAN WASSENHOVE

THE DOMAIN AND INTERPRETATION OF UTILITY FUNCTIONS: AN EXPLORATION

ABSTRACT. This paper proposes an exploration of the methodology of utility functions that distinguishes *interpretation* from *representation*. While representation univocally assigns numbers to the entities of the domain of utility functions, interpretation relates these entities with empirically observable objects of choice. This allows us to make explicit the standard interpretation of utility functions which assumes that two objects have the same utility if and only if the individual is indifferent among them. We explore the underlying assumptions of such an hypothesis and propose a non-standard interpretation according to which objects of choice have a well-defined utility although individuals may vary in the way they treat these objects in a specific context. We provide examples of such a methodological approach that may explain some reversal of preferences and suggest possible mathematical formulations for further research.

KEY WORDS: Utility, representation, interpretation, preference reversal

1. INTRODUCTION

‘The standard rationality hypothesis is that behavior can be represented as the maximization of a suitably restricted utility function’ (Barbera et al., 1998, p. ix). But what constitutes the domain of such a utility function?

It is often informally argued that the domain of utility functions encompasses ‘everything the individual is concerned about’. The entities over which a utility function assigns numbers are left formally undefined. As a formal model, a utility function represents abstract entities whose relations do not depend on a concrete interpretation. This is necessary for reflecting the structural properties of these entities and provides the model with its normative character.

On the other hand, when utility functions are used in practice to describe or predict an observable behavior, a concrete object of choice must be specified. This is necessary for the model to have *some* relation with empirical phenomena. While there is no normat-



Theory and Decision **51**: 329–349, 2001.

© 2002 Kluwer Academic Publishers. Printed in the Netherlands.

ive character without abstract entities, there is no empirical falsification without interpretation.

In the following section, we make explicit these two steps of the methodology of utility functions. We refer to *interpretation* as classifying empirical objects of choice into abstract classes of objects. Formally speaking, these abstract classes compose the domain of utility functions. We then refer to *representation* as assigning numbers univocally to these classes. The combination of interpretation and representation provides the methodology of utility functions with both a normative and an empirical character. This section clarifies these two steps. In particular, it separates representation from the standard interpretation of utility functions, i.e., that *two empirical objects of choice have the same utility if and only if the individual is indifferent between them*.

In Section 3, we focus on this standard interpretation and on its underlying invariance assumptions. Following Sen (1986), we distinguish between the assumption that similar empirical objects can be treated as identical (*substantive invariance*) and the assumption that identical objects lead to equivalent effects (*procedural invariance*). We propose to maintain the former while relaxing the latter, reflecting that behavior of individuals may not be reduced to the properties of the objects of choice. We present a methodological approach that does not assume procedural invariance and thus departs from the standard interpretation of utility functions. This approach is proposed to reflect that observed behavior also depends on the way individuals treat the objects of choice.

In Section 4, we propose an analogy and an example to illustrate such a methodological approach. The analogy discusses the measurement of masses with a biased measuring device whose bias is unobservable. While objects are assumed to have invariant properties (their mass), the bias of the balance is analogous to the way individuals may treat objects in a decision-making situation. The decision-making example discusses why an object may be assigned a lower price than another when evaluated in isolation, while being preferred when the two objects are directly compared. Such violations of rationality are well documented, in particular in the application of the willingness-to-pay method to measure environmental values. Our methodology explains them rationally by allowing be-

havior to depend on the way individuals treat objects of choice. In this example like in the analogy, empirical observation involves invariant properties of the objects of choice and varying properties for the process of choice.

In Section 5, we suggest two mathematical formulations of such a methodological approach. First, we suggest a modification of the von Neumann and Morgenstern framework that maintains all original axioms but modifies the first natural primitive. This preserves representation by an expected utility function but does not oblige one to maintain the standard interpretation. Second, we suggest a more general version of the issue in terms of algebraic structures. The theoretical problem whether we can measure procedural variations through their influence on observable behavior is interesting but not solved. As our work is exploratory, we open the discussion on these issues.

Finally, Section 6 discusses a relatively scattered literature on the issue of combining interpretation with representation.

2. THE DOMAIN AND INTERPRETATION OF UTILITY FUNCTIONS

In the standard approach to rational behavior, a utility function assigns numbers to the entities of choice so that an object is chosen over another if and only if it is assigned a higher number. For example, if an individual chooses a pear over an apple, then a utility function that represents this choice assigns a higher utility to the pear than to the apple.

The existence of such a function is established through a *representation theorem* which effectively constructs such an assignment procedure while characterizing it by its uniqueness properties. These theorems of existence rely on the construction of a one-to-one correspondence or isomorphism between numbers and the entities to which these numbers are assigned.

In our example, designating the pear by the letter a , the apple by the letter b , the relation of choice between them by \succsim , and, denoting the utility assigned to a by x and the utility assigned to b by y , we write

$$a \succsim b \Leftrightarrow x \geq y$$

where \geq is the natural relation ordering numbers. However, such a utility function u is not one-to-one when two objects have the same utility. If we want to reflect in the methodology that two distinct objects may have the same utility while still establishing the existence of a one-to-one correspondence, we cannot take the objects of choice as actually composing the domain of the utility function. Constructing a domain over which a utility function can be formally proved to exist as a one-to-one correspondence between this domain and numbers is part of the *interpretative* step.

The standard approach to interpretation assumes that two objects are assigned the same utility if and only if the individual is indifferent among them. For example, if an individual is indifferent between a pear and an orange, then the pear and the orange will be assigned the same utility. In this manner, the domain of a one-to-one utility function is constructed by first classifying objects of choice according to the indifference relation. All objects over which the individual is indifferent belong to the same class. When the indifference relation is assumed to be an equivalence relation, these classes are equivalence classes. Designating the orange by the letter a' , the relation of indifference by \sim , we would then write

$$a \sim a' \Leftrightarrow a \in [a] \text{ and } a' \in [a].$$

Formally, we have constructed the quotient set A/\sim that consists of disjoint subsets of A . These subsets are totally ordered by the relation \geq over the quotient set A/\sim . The standard approach to interpretation is summarized by

$$a \succsim b \Leftrightarrow [a] \geq [b].$$

This notation distinguishes \geq over A/\sim from \succsim over A to keep in mind that the former is antisymmetric: $[a] \geq [a']$ and $[a'] \geq [a] \Rightarrow [a] = [a']$, while the latter is not: $a \succsim a'$ and $a' \succsim a \Rightarrow a \sim a' \nRightarrow a = a'$. In other words, there is no possibility to distinguish between two classes that have the same utility, since these two classes are one and only one identical class. This is necessary to establish the existence of a one-to-one correspondence between equivalence classes and numbers, the *representation* step.

Representation establishes the existence of a one-to-one correspondence between equivalence classes and numbers. Because inter-

pretation has already regrouped all objects of choice that are mutually indifferent into one single class, there is no possibility for the same number to be assigned to two classes and the utility function from equivalence classes to numbers is one-to-one. Representation can be summarized by

$$[a] \geq [b] \Leftrightarrow x \geq y$$

with $u([a]) = x$ and $u([b]) = y$. Formally speaking, the domain of the utility function u does not consist of the objects of choice but of abstract equivalence classes. In our example, the utility x assigned to the pear designated by a is not the utility of a but the utility of an abstract entity to which a is supposed to belong and that we denoted by $[a]$.

There are therefore two distinct steps in the methodology of utility functions. First, regrouping all objects of choice into abstract equivalence classes and second, assigning univocally one number to each of these equivalence classes. In this manner, the methodology of utility functions combines interpretation and representation.

The disentanglement of interpretation from representation leads to express the methodological statement that an object is chosen over another if and only if it is assigned a higher utility as being implied by two distinct equivalences, an interpretive statement and a formal statement

$$a \succ b \Leftrightarrow [a] \geq [b] \quad \text{interpretative statement}$$

and

$$[a] \geq [b] \Leftrightarrow x \geq y. \quad \text{formal statement}$$

imply

$$a \succ b \Leftrightarrow x \geq y. \quad \text{methodological statement}$$

These two steps are shown in Figure 1.

Consider a situation in which the apple is chosen over the orange. We have $b \succ a'$ while $x > y$ contradicting the methodological statement $x \geq y \Leftrightarrow a' \succ b$. The question is which one of the formal statement $x \geq y \Leftrightarrow [a'] \geq [b]$ or the interpretive statement $[a'] \geq [b] \Leftrightarrow a' \succ b$ is violated?

A standard approach would assume the interpretative statement and reject the formal statement. In other words, the pear, the apple




Figure 1. Combining interpretation with representation.

and the orange cannot be assigned the utility x , y and x respectively. In particular, the pear and the orange *cannot* belong to the same equivalence class.

What if one assumes the formal statement and rejects the interpretive statement? This would mean that the pear and the orange would belong to the same equivalence class and would be assigned the same utility but that the individual may not treat them the same way in the two situations. Disentangling interpretation from representation thus allows one to consider the possibility that empirical falsification of the methodology of utility function does not stem from an improper isomorphism between the domain of the utility function and numbers but from an improper interpretation of what utility means and relates to empirical observation.

How to express such a possibility for the influence of interpretation on choice without falling into arbitrariness requires to specify distinct underlying assumptions of interpretation that have remained implicit throughout this section.

3. INTERPRETATION WITHOUT PROCEDURAL INVARIANCE

Considering that abstract entities can be substituted for empirical objects of choice without affecting the representation of choice is the standard approach to interpretation. It considers that there is no qualitative distinction between an empirical object of choice and the abstract equivalence class to which it belongs. But there must be some distinction since two objects of choice that belong to the same class can be distinguishable as objects of choice. Therefore, each object of choice possesses properties that are not properties of the class to which it belongs. The standard approach to interpretation

merely assumes that these properties do not have any empirical influence. In mathematics, the process by which the quotient set A/\sim is constructed is called an *abstraction of qualities*. Such an abstraction occurs precisely because the relation between an object of choice and its abstract entity is not one-to-one. This is another manner to formulate the distinction between the process of interpretation and the process of representation as one implies a loss of information while the other reflects all available information in a richer structure.

The standard interpretation consists in a reduction of choice to well-defined objects of choice so that choice does not depend on the situation in which the act of choice is performed. In our example, the individual who is indifferent between a pear and an orange is supposed to remain indifferent among a similar pear and a similar orange in a similar situation of choice. First, it is assumed that the objects of choice are well-defined and that they remain the same. The orange compared with the pear is assumed to be identical to the orange compared with the apple although the two oranges do appear in distinct situations of choice. This assumption about the invariance of the objects of choice $a, b, a' \dots$ is called *substantive invariance*. Second, it is assumed that the individual who is indifferent between the objects also remains the same. In other words, the pear and the orange are interpreted in the same manner across situations of choice. This assumption about the invariance of the indifference relation \sim is called *procedural invariance* because the two objects are assumed to be treated in the same way. The standard interpretation thus assumes that similar objects across contexts are identical (substantive invariance) and that identical objects have similar effects across situations (procedural invariance), i.e. that similar objects of choice have similar effects.

However, similar objects across situations may not be identical. A typical approach consists in indexing objects by the situation in which they appear, thus defining new objects of choice. Different proposals have been made in this respect which are well developed and sometimes constitute avenues of research on their own. A main issue of such a relaxation of substantive invariance is however to compare these newly defined objects. If objects are distinct across situations and observation takes place in a given situation, it becomes problematic to formulate observable primitives across situ-

ations. When relaxing substantive invariance, the challenge is to maintain a predictive character for the methodology at the level of individual choice.

Departing from such an approach, we explore a relaxation of procedural invariance. We assume that similar objects across contexts are identical but that they may be treated in different manners by the individual in distinct situations of choice. Such an approach differs from the preceding one in the sense that objects of choice remain well-defined across situations since they are assumed to be identical (substantive invariance). It is the intensity of their relation to the individual who chooses which is not assumed to remain invariant. In our example, the pear, the apple and the orange are assumed to remain the same as objects of choice in the different situations while the individual is allowed to treat the same orange differently, for instance, depending on whether it is compared to a pear or to an apple.

Because each object of choice is well-defined across contexts (substantive invariance), each object belongs to one and only one equivalence class. This class retains only some of the invariant properties of the object and does not depend on the process by which an object is treated in a specific situation of choice. It is therefore possible to prove, given adequate properties for these equivalence classes and the relations among them, the existence and uniqueness of a utility function that represents a total ordering of these classes. In this manner, the utility of an object of choice a is interpreted as representing only those properties of a that are independent of the situation. We could say that *the utility of an object is an absolute measure of its substantive properties*.

Since relaxing procedural invariance amounts to admitting that an object may be treated differently across contexts, a given object may lead to different empirical effects depending on the situation. In particular, choice may depend on the other objects to which it is compared. Besides the properties of the object represented by its utility, choice is influenced by *the process by which the individual interprets an object relatively to others* in a specific situation. In a context influencing more a than b , an object a of utility x may be chosen over an object b of greater utility $y \geq x$. Although the pear, the apple and the orange are supposed to have a well-defined utility



Figure 2. Choice in favor of a .

as objects of choice, we do not suppose that, for instance, the orange is treated in the same way when compared with the pear and when compared with the apple.

In this relaxation of the standard interpretation, we disentangle the properties of the objects from the properties of the individual, as a subject who interprets the object and actually carries out the act of choice. Such a relaxation introduces more flexibility but only at the interpretive level. Because substantive invariance is maintained, there remains the same structure at the formal level than with the standard interpretation. Whether such an approach allows one to better understand observation of empirical phenomena is now explored through two examples.

4. OBSERVABLE BEHAVIOR WITHOUT PROCEDURAL INVARIANCE

The distinction between the properties of the objects of choice and the properties of the subject who carries out choice itself can be illustrated by considering a black box out of which we can only observe the empirical effects of choice. This behavioral approach assumes that the inner properties of the individual are not directly observable. It is also a standard approach in natural sciences, a typical example being the measurement of mass. This section illustrates the observation of behavior when procedural invariance is not assumed, first for the measurement of mass and then in decision-making.¹

We consider a two-arm-balance hidden in a black box. With such a balance, we observe that an object a is 'chosen' over another object b although we do not observe how the objects are 'treated' by the black box (Figure 2).

Independently of how each object is treated by the balance, it has a well-defined mass and there is an isomorphism between masses and numbers. However, which object is chosen by the balance does not only depend on the masses of the objects but also on how these




Figure 3. b has a higher mass than a .

objects are ‘treated’ by the balance. Because the balance is not assumed to have arms of equal length, an object with a lower mass may be chosen over another object of greater mass. Naturally, the definition of mass can no longer be expressed as the property of objects for which the balance is at equilibrium. Such a definition implicitly assumes that each object on the balance is treated in identical ways: the very assumption we want to relax.

The definition of masses as equivalence classes of objects whose properties are similar can nevertheless be carried out by an operation of *substitution* of one object for another. Such an operation does not involve the two objects placed on the balance but only one of them, say a , and another one, say a' , which is not on the balance but which, if placed in substitution of a , would lead to the same effect than a . In this manner, the class $[a]$ of objects a, a' regroups all objects of the same mass although two objects of the same mass may not balance the measuring device because of its bias. Figure 4 illustrates an object a chosen over an object b although the mass of b is greater than the mass of a .

In this example, the influence of the context on the measuring device is reflected by the position of the fulcrum of the balance. The behavior of the biased balance cannot be reduced to the mass of the objects that are placed on it but also depends on the balance itself. Empirical observation, together with the assumption that objects of choice have well-defined properties, allows us to appreciate the position of the fulcrum of the balance. How and to which extent we can properly measure such a position on the sole basis of the measured masses is an interesting question on which we come back in the next section. We now illustrate behavior without procedural invariance with a decision-making example.

In decision making, an interesting class of phenomena violating the standard interpretation of utility functions occurs when an individual is observed to choose between two objects and then asked



Figure 4. Choice in favor of a .

to evaluate each object separately, for instance by assigning them a price. In some systematic cases, rational individuals are observed choosing an object a over an object b while assigning a higher price to b than to a .² Observing an individual choosing an object a over an object b can be illustrated with the same figure than for the measurement of mass. The individual, as a measuring device, is hidden in the black box and his inner values are not accessible to direct empirical observation (Figures 4 and 5).

As for the measurement of mass, relaxation of procedural invariance necessitates to distinguish between empirical comparison of two objects of choice and abstract substitution of one class for the object itself. Observing an individual assigning a price to object a can then be interpreted as a substitution between the object a and the class to which a belongs. Being absolute, the operation of substitution thus does not involve the other object. Being relative, choosing one object over another involves both objects and can be influenced by the propensity of the individual to favor one object over the other.

A typical type of situation when such reversals occur is when we attempt to evaluate environmental concerns through pricing, i.e., willingness to pay techniques. A seminal example being that most individuals confronted with the choice between improving the air quality in their town and adding a VCR to their TV prefer to improve the air quality while they assign a higher price to the addition of the VCR to their TV (Irwin et al., 1993).

In the interpretation explored here, this would mean that the value of environmental concerns must not be reduced to the value of the object we desire to protect but also integrates a procedural value singular to the individual who chooses. Therefore, reducing environmental concerns to their willingness-to-pay values would underestimate individuals motivation to act in order to preserve their environment.




Figure 5. b has a higher price than a .

In this example, pricing differs from choice and the distinction corresponds to the one between substitution and comparison. In such an interpretation, there is no irreducible contradiction in the behavior of the individual who assigns a higher price to an object that is not chosen. Naturally, absolute evaluation and relative comparison are equivalent under the assumption that choice is not influenced by the process of choice (procedural invariance). In the general case, observing a discrepancy between absolute evaluation and relative comparison informs us about how the objects of choice are treated by the individual, i.e., about *the process of choice*. Like for the measurement of the arms of the balance, the question remains to which extent can we measure the influence of the process of choice.

5. MATHEMATICAL SUGGESTIONS

A mathematical characterization of the methodology of utility functions without assuming procedural invariance remains in its exploratory phase. In this respect, we first suggest a formulation based on the axiomatization of expected utility theory proposed by von Neumann and Morgenstern in *Games and Economic Behavior* (1953).

The treatment of the methodology of utility functions proposed by von Neumann and Morgenstern carefully distinguishes the formal and the interpretive steps. The mathematical treatment is strictly restricted to the formal step of representation. A standard interpretation is presented conceptually and various remarks along the text warn that such a step is not part of the formal treatment. Typically, the authors do not speak of a utility function per se but of a correspondence between some abstract entities called '*utilities*' and numbers. As pointed out by their following commentators (Marschak, 1950; Malinvaud, 1952, see also Fishburn, 1989, and Fishburn and Wakker, 1995), such utilities are equivalence classes, like those de-

noted above by $[a]$, $[b]$, Axioms apply to these classes and a theorem demonstrates the existence and the uniqueness properties of the correspondence between them and numbers. Such a formulation helps us to suggest a variation of the von Neumann and Morgenstern framework that differs only by its interpretation, not by its formal axioms.

The standard interpretation of the von Neumann and Morgenstern axioms is expressed through one ‘natural relation’ and one ‘natural operation’ among utilities. The first natural relation is denoted by $[a] > [b]$ and reads ‘ $[a]$ is preferable to $[b]$.’ The second natural operation introduces a number λ strictly between 0 and 1 and is denoted by ‘ $\lambda[a] + (1 - \lambda)[b] = [c]$ ’. It reads ‘utilities can be combined with probabilities’. The implicit interpretation is that a is chosen over b if and only if $[a]$ is preferable to $[b]$, that is $a \succ b \Leftrightarrow [a] > [b]$, and that a is indifferent to b if and only if they have the same utility, that is $a \sim b \Leftrightarrow [a] = [b]$. Again, such an interpretation is not part of their formal model.

In order to reflect the influence of procedural variations of interpretation, we suggest to modify the first natural relation while maintaining the second natural operation, the latter corresponding to substantive invariance. The modification assumes that empirical observation of choice reveals *a relation among weighted equivalence classes*. The natural primitives of the framework become a system of equivalence classes $[a]$, $[b]$, $[c]$ for which we consider:

$$\begin{aligned} \text{First Natural Relation : } & \alpha[a] > (1 - \alpha)[b], \\ \text{Second Natural Operation : } & \lambda[a] + (1 - \lambda)[b] = [c], \\ & 0 < \alpha < 1 \text{ and } 0 < \lambda < 1. \end{aligned}$$

The first natural relation can be read as ‘ a is chosen over b ’. The unknown number α is a ‘hidden bias’ that is not given a priori. The known number λ keeps its interpretation as a probability that is objectively given. The framework allows one to express that ‘ a is chosen over b ’ while ‘ $[b]$ is preferable to $[a]$ ’ suggesting a bound $\alpha > \frac{1}{2}$ revealed a posteriori from empirical observation of choice. We can indeed keep all the axioms of von Neumann and Morgenstern combining utilities and the relation among them having abstracted the process by which these utilities are interpreted, i.e. the bias α . As they were expressed by their authors, these axioms are:

- (1) One and only one of the following holds: $[a] > [b]$, $[b] > [a]$, $[a] = [b]$;
- (2) Transitivity: $[a] > [b]$ and $[b] > [c]$ imply $[a] > [c]$;
- (3) Monotonicity: $[a] > [b]$ implies $[a] > \lambda[a] + (1 - \lambda)[b]$;
- (4) Continuity: $[a] > [c] > [b]$ implies $\lambda[a] + (1 - \lambda)[b] > [c]$ for some λ ;
- (5) Commutativity: $\lambda[a] + (1 - \lambda)[b] = (1 - \lambda)[b] + \lambda[a]$;
- (6) Associativity: $\mu(\lambda[a] + (1 - \lambda)[b]) + (1 - \mu)[b] = \mu\lambda[a] + (1 - \mu\lambda)[b]$.

Although the presence of these axioms ensures the existence of a univocal correspondence with numbers, the uniqueness of such a correspondence is restricted to the group of positive affine transformations. This means in particular that the structure does not distinguish between different representations that vary with their origin or, equivalently, that we cannot compare ratios of utilities. As a consequence, even a posteriori, a bias α can solely be shown to be less or greater than $\frac{1}{2}$, without being quantitatively measured. If utilities can only be measured on an interval scale, biases can only be ranked on an ordinal scale.

The formulation of the bias reflecting procedural variations of interpretation as weights on each side of the relation thus suggests a form of indeterminacy (think of comparing 1 unit of the first rank with 2 units of the second rank). Whether such a combination of a quantitative scale (here interval) with a qualitative scale (here ordinal) together with the partial indeterminacy that results are necessary features of a mathematical formulation that combines procedurally biased interpretation with representation remains to be clarified.³

An interesting feature of the von Neumann and Morgenstern formulation is that it does not use any axiom of independence. As Malinvaud (1952) explains, such an axiom is part of the standard interpretation and has no place per se at the abstract level of equivalence classes (see also Fishburn and Wakker, 1995). On the other hand, an important effort of research has been directed to independence because it appears the most commonly violated axiom of expected utility theory. Consequently, violations of independence may be a violation of the standard interpretation, not a violation of the axioms of von Neumann and Morgenstern.

The idea we want to reflect is that independence holds at the level of equivalence classes while not at the level of the objects of choice. To formulate such an influence of interpretation between the abstract structure of equivalence classes and the empirical structure of objects of choice, we rely on the definition of an extensive structure (Krantz et al., 1971). An extensive structure is a quotient set endowed with a binary operation \oplus and a total ordering $>$. An extensive structure satisfies the following properties.

- (1) One and only one of the following holds: $[a] > [b]$, $[b] > [a]$, $[a] = [b]$;
- (2) Transitivity: $[a] > [b]$ and $[b] > [c]$ imply $[a] > [c]$;
- (3) Monotonicity (independence): $[a] \geq [b]$ iff $[a] \oplus [c] \geq [b] \oplus [c]$;
- (4) Solvability: $[a] > [b]$ implies $[a] = [b] \oplus [c]$ for some c ;
- (5) Archimedean: $n[a] > [b]$ for some integer n , with $1[a] = [a]$ and $n[a] = (n - 1)[a] \oplus [a]$;
- (6) Commutativity: $[a] \oplus [b] = [b] \oplus [a]$;
- (7) Associativity: $[a] \oplus ([b] \oplus [c]) = ([a] \oplus [b]) \oplus [c]$.

When we assume the standard interpretation, these axioms are expressed in terms of objects of choice a, b, c, \dots , the weak ordering relation \succsim is substituted for the total ordering \geq and the operation \circ among objects of A is substituted for the operation \oplus among equivalence classes of A/\sim . Using the standard interpretation to define the ordering among equivalence classes as $a \succsim b \Leftrightarrow [a] \geq [b]$, we obtain the axioms of an extensive structure.

Our suggestion is thus to impose axioms on the objects of A sufficient to construct an equivalence relation \approx , in general distinct from indifference \sim , on A compatible with \circ so that the quotient set A/\approx is an extensive structure totally ordered by a relation \geq . We want however to relax the independence condition. We propose the following sets of axioms for a 'semi-monotonic structure'.

- (1) One and only one of the following holds: $a \succ b$, $b \succ a$, $a \sim b$;
- (2) Transitivity of \succ : $a \succ b$ and $b \succ c$ imply $a \succ c$;
- (3) Semi-monotonicity: $a \succsim b$ implies $a \circ c \succsim b$; $a \succsim b \circ c$ implies $a \succsim b$ for all c .
- (4) Solvability: $a \succ b$ implies $a \succsim b \circ c$ for some c ;

- (5) Archimedean: $na \succ b$ for some integer n , with $1a = a$ and $na = (n - 1)a \circ a$;
- (6) Commutativity: $a \circ b = b \circ a$;
- (7) Associativity: $a \circ (b \circ c) = (a \circ b) \circ c$.

With such a structure, we may construct an equivalence relation \approx on A so that there exist a total ordering $>$ and an operation \oplus on A/\approx so that $\langle A/\approx, >, \oplus \rangle$ is an extensive structure.⁴ The semi-monotonic structures $\langle A, \succ, \circ \rangle$ indeed compose a family of structures, one of which reflecting the standard interpretation. In this special case, we do have $a \sim b \Leftrightarrow [a] = [b]$ and the two relations \sim and \approx are identical (with $a \sim b \Leftrightarrow a \not\sim b$ and $b \not\sim a$). In general, the indifference relation \sim is not necessarily transitive and differs from the equivalence relation \approx on A . The question is whether we can order these semi-monotonic structures so as to measure how much a given indifference relation \sim departs from the equivalence relation \approx . This amounts to order the empirical relations \succ according to their departure from the standard interpretation. In the case of the biased balance, the question is whether this allows one to properly measure the bias of the balance from the objects that are measured by it.

6. DISCUSSION OF THE LITERATURE

The literature about interpretation and its relation to formal models of utility functions is relatively scattered. The standard interpretation that two objects have the same utility whenever the individual is indifferent among them seems to date back to the 19th century. A historical approach to the notion of utility can be found in Stigler (1950). For utility functions, the distinction between interpretation and representation seems to date back to the first formal treatment proposed by von Neumann and Morgenstern (1947). At that time, it had already generated considerable controversy about the nature of the domain over which a utility function is defined (see, e.g., Marschak, 1950; Malinvaud, 1952). With Herstein and Milnor (1953), interpretation is again subsumed in the exposition of the expected utility representation theorem, an approach that we can find in the more general and modern theory of representational

measurement (Krantz et al. 1971). This is the formulation of the standard interpretation we have chosen to expose.

Interpretation, as the relation between empirically observable variables and their formalization as abstract variables has been acknowledged as a fundamental assumption of measurement. It is indeed an assumption that is 'ubiquitous throughout all science' (Luce and Narens, 1986; see also Luce, 1996). That primitive entities of formal models are hypothetical and thus not the actual empirical objects of choice remains nevertheless subject to controversy. A recent clarification, due to Aumann, is however explicit :

'In any axiomatic system, the arguments depend crucially on hypothetical, artificial situations that never existed. The essence of the axiomatic approach is that it works with an entire system, and with the relations between the objects in it. In particular, it relates the given, 'real', situation to a whole lot of other situations, all of them hypothetical. This happens constantly – in Arrow's (1951) social welfare theory, in the Shapley value (1953), in Nash's (1950) bargaining solution, indeed even in Savage's (1954) development of probabilities themselves!' (Aumann, 1998, p. 935).

The use of classes as primitives is discussed in de Finetti, in his seminal paper about subjective probabilities (1937) and later in the two volumes of the *Theory of Probability* (1974).⁵

An important and specific discussion of the role of the standard interpretation can be found in Sen (1986). The distinction between substantive and procedural invariance is explicitly discussed, as well as the importance of such forms of invariance for the independence condition in expected utility theory, but also with regards to Arrow's possibility theorem. The argument is stated in terms of information structure but this is equivalent. Interpretation is seen as drawing a line between relevant information and discarded information:

'Any principle of choice uses certain types of information and ignores others. A principle can be understood and assessed in terms of the information that it demands and the information it rules out.' (ibid., p. 29)

Another work where the issue of varying interpretations appears is in (Sounderpandian, 1992). In this multi-criteria approach, interpretation is seen as the choice of a particular criterion space. The author states that the order-preserving character of the correspondence between empirical objects and their multi-criteria interpretation is a matter of 'belief', not of an axiomatic approach that is empirically falsifiable.

The idea that process preferences may force one to depart from the standard methodology inherited from natural sciences is evoked in Sen (1997). The idea of maintaining some substantive invariance while relaxing procedural invariance is related to the notion of process preferences and process utility developed in Le Menestrel (1999, see also 2001 and references). These works introduce and model considerations for the processes in order to reflect truly subjective concerns of the individual who acts in a particular context. It is exploratory in nature since we have been unable to find other rigorous treatments of specific process considerations in the literature.

The formulation of the postulate $\alpha u > (1 - \alpha)v$ as a modification of von Neumann and Morgenstern first natural relation was first proposed in Le Menestrel (1998). The partial indeterminacy that results from such a combination of an ordinal and a cardinal measure may seem awkward at first, but not so in light of the search for models that are more open than formal models and may better reflect the nature of rational behavior. In this respect, the combination of the different types of hypotheses to give rise to empirical observation is discussed in Starmer (1999) and outside utility theory, for instance in Quine (1975).

Unless we have missed such a work, it remains however to formulate a full characterization of the system formed by representation and interpretation without procedural invariance. In this respect, we are following the idea of relaxing the transitivity of indifference in what we proposed to call semi-monotonic structures. A review of intransitive indifference can be found in Fishburn (1970). We are well aware that we do not provide such a characterization but merely suggest some possible directions for research. We wish to see such a construction in the future.

7. CONCLUSION

This paper proposes a methodological exploration of utility functions that attempts to clarify the distinct roles of representation and of interpretation. While the former relies on theorems proved from axioms, the latter is a matter of definition. For a methodology that claims both a normative character and some relation with empirical

observation, the domain of utility function does not consist of the empirical objects of choice but of abstract entities that are related to them through interpretation. The natural way to carry out such an interpretation is isolated and studied as the standard interpretation of utility functions.

We suggest that such a standard interpretation of utility functions, inherited from the scientific methodology primarily developed for the natural sciences, may not properly consider the inner properties of the individual who chooses. This may be because scientific methodology has been primarily interested in the studies of the properties of objects independently from the process by which these properties are observable. On the other hand, individuals who carry out the act of choice are not invariant in the way they treat objects and their behavior may thus depend on the context of choice beyond the properties of empirical objects of choice. In other words, empirical violations of the methodology of utility functions may also be explained by a variance of the properties that are not reflected in the standard interpretation. Rather than looking for the 'true laws of behavior', we suggest to better understand the process by which laws are contextualized and give rise to empirical observation of behavior.

Our exploration suggests possibilities to reflect the influence of the context of choice in models that do not rely on the standard interpretation. We propose examples of such models and suggest some mathematical treatment. Further research should better relate such models with current methodological approaches so as to better assess their potential interest.

NOTES

1. The parallel between the measurement of mass and the measurement of utility is discussed in von Neumann and Morgenstern (1953) and is indeed a natural one as the balance is an ancestral analogy for measurement (see Krantz et al., 1971). One can find a discussion of the black box with regards to the axiomatic method in von Neumann (1951, p. 2).
2. Hsee et al. (1999) review these phenomena and propose a psychological analysis of these types of preference reversals between joint and separate evaluation.

3. An application of such a framework to the Allais paradox and the utility for gambling can be found in Le Menestrel (2001). An application to procedural concerns in game theory and a more general discussion can be found in Le Menestrel (1999).
4. Work on the characterization of the biased balance with such structures is currently carried out with Bertrand Lemaire.
5. See in particular Chapter 6 (1937), Chapter 2 (1974) in particular remark 2.4.3 and appendices 13 and 17.

REFERENCES

- Arrow, K.J. (1951), *Social Choice and Individual Values*, New York: John Wiley and Sons.
- Aumann, R.J. (1998), Common priors: A reply to Gul, *Econometrica* 66, 929–938.
- Barbera, S., Hammond, P.J. and Seidl, C. (eds) (1998), *Handbook of Utility Theory: Vol. 1: Principles*. Dordrecht / Boston / London: Kluwer Academic Publishers.
- De Finetti, B. (1937), *La prévision: ses lois logiques, ses sources subjectives*. Paris, France: Institut Henri Poincaré.
- De Finetti, B. (1974), *Theory of Probability*, Chichester, UK: John Wiley and Sons.
- Fishburn, P.C. (1970), Intransitive indifference in preference theory: A survey, *Operations Research* 18, 207–228.
- Fishburn, P.C. (1989), Retrospective on the utility theory of von Neumann and Morgenstern, *Journal of Risk and Uncertainty* 2, 127–158.
- Fishburn, P.C. and Wakker, P. (1995), The invention of the independence condition for preferences, *Management Science* 41, 1130–1144.
- Herstein, I.N. and Milnor, J. (1953), An axiomatic approach to measurable utility, *Econometrica* 21, 291–297.
- Hsee, C.J., Loewenstein, G.F., Blount, S. and Bazerman, M.H. (1999), Preference reversals between joint and separate evaluation of options: A review and theoretical analysis, *Psychological Bulletin* 125, 576–590.
- Irwin, J.R., Slovic, P., Lichtenstein, S. and McClelland, G.H. (1993), Preference reversals and the measurement of environmental values, *Journal of Risk and Uncertainty* 6, 5–18.
- Krantz, D., Luce, R.D., Suppes, P. and Tversky, A. (1971), *Foundations of Measurement, Volume 1: Additive and Polynomial Representations*. New York and London: Academic Press.
- Le Menestrel, M. (1998), A note on embedding von Neumann and Morgenstern utility theory in a qualitative context, INSEAD Working Papers, 52.
- Le Menestrel, M. (1999), A model of rational behavior combining processes and consequences. Unpublished Ph.D. Dissertation.

- Le Menestrel, M. (2001), A process approach to the utility for gambling, *Theory and Decision* 50, 249–262.
- Luce, D.R. (1996), The ongoing dialog between empirical science and measurement theory, *Journal of Risk and Uncertainty* 5, 5–27.
- Luce, D.R. and Narens, L. (1987), Measurement scales on the continuum, *Science* 236, 1527–1532.
- Malinvaud, E. (1952), Note on von Neumann-Morgenstern's strong independence axiom, *Econometrica* 20, 679.
- Marschak, J. (1950), Rational behavior, uncertain prospects, and measurable utility, *Econometrica* 18, 111–141.
- Nash, J.F. (1950), The bargaining problem, *Econometrica* 18, 155–162.
- Quine, W.V.O. (1975), On empirically equivalent systems of the world, *Erkenntnis* 9, 313–328.
- Savage, L.J. (1954), *The Foundations of Statistics*, New York; 2nd edition revised and enlarged 1972, Dover.
- Sen, A. (1986), Information and invariance in normative choice. In *Social Choice and Public Decision Making* (pp. 29–55). Cambridge, UK: Cambridge University Press.
- Sen, A. (1997), Maximization and the act of choice, *Econometrica* 65, 745–779.
- Shapley, L.S. (1953), A value for n-person games. In: H. Kuhn and A.W. Tucker, (eds.), *Contributions to the Theory of Games*, Vol. II (pp. 305–317). Princeton, NJ, Princeton University Press.
- Sounderpandian, J. (1992), Transforming continuous utility into additive utility using Kolmogorov's theorem, *Journal of Multi-Criteria Decision Analysis* 1, 93–99.
- Starmer, C. (1999), Experiments in economics: Should we trust the dismal scientists in white coats?, *Journal of Economic Methodology* 6, 1–30.
- Stigler, G. J. (1950), The development of utility theory: I, II, *Journal of Political Economy* 58, 307–327, 373–396.
- von Neumann, J. and Morgenstern, O. (1944), *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press. Second edition 1947, Third edition 1953.
- von Neumann, J. (1951), The general and logical theory of automata. In: L.A. Jeffries (ed.), *Cerebral Mechanisms in Behavior: The Hixon Symposium*. New York: John Wiley and Sons.

Addresses for correspondence: Marc Le Menestrel, Universitat Pompeu Fabra, Departament d'Economia i Empresa, Ramon Trias Vargas, 25-27, E-08005, Barcelona, Spain

E-mail: marc.lemenestrel@econ.upf.es