Informing Science: the International Journal of an Emerging Transdiscipline **Volume 12, 2009**

A Comment on 'A Psychologically Plausible Goal-Based Utility Function'

Robert Bordley
General Motors & University of Michigan, Ann Arbor, MI, USA

Robert.bordley@gm.com

Abstract

Gill advocated using goal-based utility functions versus the standard commodity-based neoclassical utility function. He argued that a goal-based utility function would resolve many observed inconsistencies between what economics predicts and the behavior observed in psychological experiments. To support Gill's argument, this paper focuses on the standard extension of neoclassical utility to gambles. Given this extension, it is straightforward to rewrite the neoclassical utility function as a goal-based utility function.

This goal-based utility function will correspond to quantities like the Sharpe ratio and effect size already widely used in finance and medicine. As we show, minor adjustments of this formula are mathematically identical to proposed psychological plausible models of human behavior. Finally we show how this formulation can be extended to the case of a goal with multiple sub-goals.

Keywords: Goal, Utility, Economics, Effect Size, Sharpe Ratio, Decision Analysis

Introduction

An individual's neoclassical utility function is typically defined over their current situation (e.g., how much of various commodities they own) and is not explicitly a function of the individual's needs and aspirations. Thus it does not typically describe the executive whose utility depends on how much his income exceeds the income of his peers. For these and other reasons, Gill (2008) proposed replacing this neoclassical utility with a utility that explicitly depends on the individual's goals. Gill's seminal paper showed that this goal-based utility function addresses many violations of utility theory observed in psychological experiments. He also discussed how this utility function, unlike the stable neoclassical utility function, could change as the individual learned.

This paper shows how Gill's goal-based utility function can be explicitly written in terms of the neoclassical utility function. We show that this goal-based utility also appears (under other names) in different fields. We then show that small changes in this goal-based utility leads to models that can explain some of the cases in which the neoclassical utility function is often inconsistent with observed human behavior. The paper closes by discussing how goal-based utility includes those cases in which an individual uses multiple criteria in defining his or her utility.

Editor: T. Grandon Gill

The Goal-Based Utility

The neoclassical utility function is defined as a function of known quantities (e.g., known levels of consumption of various goods.) But in many cases, individuals often make purchasing decisions about items which do not yet have known consequences. For example, an individual may own a lottery ticket which has some chance of giving the individual $1000 and some chance of giving the individual nothing. But even an individual who owns a car is typically uncertain about how many repairs the vehicle may require in the future. Thus many items should be viewed as gambles.

Since some items are gambles, the neoclassical utility function must be written in terms of gambles (as opposed to known quantities.). To rewrite the utility function appropriately, suppose gamble i leads to a known outcome x with probability p(x|i). (The outcome x could be a vector specifying the individual's consumption of many different products.) Based on plausible assumptions about individual rationality, Von Neumann and Morgenstern (1944) showed that the utility of this gamble, U(i), can be written as

$$U(i) \;=\; \sum_x U(x)\, p(x|i)$$

where U(x) is a strictly monotonic function of the neoclassical utility function and represents the utility of a gamble which always yields the same outcome. In an extension of Borch (1968), Castagnoli and LiCalzi (1996) and Bordley and LiCalzi (2000) showed that for any utility function, U, there always exists a benchmark gamble 0 such that U(i) is just the probability of the individual considering the outcome of gamble i to be no worse than the outcome of gamble 0.

For example, an individual may want a vehicle that will meet his future needs. But the individual—who intends to keep the vehicle for five years—may not know the maximum number of people he or she will need to carry comfortably in the car or the maximum amount of luggage he or she will need to carry in the car's trunk. (In other words, the individual is uncertain about what is exactly required to meet his or her needs.) In this case, the uncertainty in gamble 0 reflects the individual's uncertainty about the future demands which will be placed on the vehicle. Suppose the individual's goal is to have a vehicle which achieves an outcome (i.e., offers a capability) that is no worse than the outcome of gamble 0 (i.e., which is enough to meet all his or her future needs.) Then U(i) is just the probability of the individual successfully achieving that goal with vehicle i.

As another example, suppose our investment goal in 2012 is to earn at least as much money as we would have earned if we had invested our money in an index fund (whose return matches the 2012 return of the S&P 500). Then gamble 0 will be the uncertain return from the index fund and U(i) will be the probability that fund i earns at least as much money as the index fund. (Under conventional statistical assumptions about stock market prices, U(i) will be identical to the Sharpe Ratio, a widely used financial measure of the relative performance of fund i.)

As a final example, suppose that medical patients in an experimental group are exposed to some new experimental treatment while individuals in a control group are exposed to a more conventional treatment. Let gamble 0 be the state of health of an individual randomly drawn from the control group and let gamble i be the state of health of an individual randomly drawn from the experimental group. Then U(i) is the probability that an individual from the experimental group is at least as healthy as an individual in the control group. Given this interpretation, U(i) is identical (Bordley, 2009) to effect size measures widely used to describe the efficacy of clinical treatment i.

Psychological Implications

This theoretical equivalence between the utility function and the probability of outperforming a benchmark gamble provides an alternate representation of utility theory which, with changes in minor auxiliary assumptions, leads to theories explaining some violations of utility theory.

For example, Oden and Lopes (1999)'s psychological theory of choice assumed that individuals made choices in light of a 'fuzzy' aspiration level and a 'fuzzy' security level and, in their view, had the same explanatory power as Kahneman and Tversky's Nobel-Prize winning cumulative prospect theory. In their theory, the 'fuzzy' aspiration level means that the individual's aspiration level is not precisely definable. In contrast, goal-based utility assumes that the aspiration level is definable but the individual is uncertain about the specific requirements that define achievement of that aspiration level. For example, an individual may aspire to be a 'success' without being certain about what is required to be considered a success. Thus the economic goal-based utility becomes Oden and Lopes theory of 'fuzzy goal-oriented' behavior by replacing a goal with uncertain requirements by a fuzzy goal.

To ensure consistency with conventional utility theory, we must assume that the outcomes of gamble 0 and gamble i are independent. This implies, for example, that our uncertainty about the capabilities of a vehicle can be assumed independent of our uncertainty about the kinds of tasks we will need to have the vehicle perform. But a generalization of conventional utility theory, state-dependent utility theory (Schervish, Seidenfeld, & Kadane,1990) , assumes that utility varies with the situation of the individual (i.e., a suffering patient may prefer a shorter life to a longer life while another individual might prefer a longer life to a shorter life). This will be consistent with goal-based utility function if we relax the assumption of the outcomes of gambles i and 0 being independent. For example, executives experience uncertainty about their own income and about the income of their peers that will typically be correlated, since everyone's income will usually rise with favorable economic times.

Hence some violations of utility theory are interpretable, from the perspective of goal-based utility, as violations of secondary assumptions used in goal-based utility. But since the neoclassical representation of utility theory does not explicitly make use of these secondary assumptions, explaining violations of neoclassical utility theory is considerably more difficult.

Extension to Multiple Goals

In many cases, the utility of some choice depends on many different attributes of the choice. Thus foods are typically evaluated on the basis of such criteria as their taste and the amount of various nutrients they supply. As Bordley and Kirkwood (2004) noted, standard assumptions in multiattribute utility theory allow us to define the goal-oriented utility in terms of goals on each of the various criteria. Specifically

(1) Define benchmarks on each attribute (e.g., benchmark gambles for taste, amount of vitamin A, amount of vitamin B, etc.).

(2) Define an outcome as 'successful on an attribute' if it generated an outcome no worse than the outcomes generated by the benchmark gamble on that attribute.

(3) Define an attribute set, A, as a collection of some (but not necessarily all) attributes. Thus A might consist of no attributes, might consist of the attribute 'taste', might consider of the attribute 'vitamin A concentration', might consist of both the attributes 'taste' and 'vitamin A concentration', etc.

(4) Define $u(i|A)$ as the probability choice i is successful on all the attributes in A and unsuccessful on all other attributes.

(5) Define w(A) as the probability that any choice will be an overall success if it was successful on the attributes in A (and unsuccessful on all other attributes).

Then simple probability theory implies

$$U(i)= \sum A\ u(i|A)\ w(A)$$

Because of its generality, this same formula also describes the probability of many complex systems not failing as a function of the different components of the complex system not failing. For example, it was the formula used for assessing the safety of nuclear power plants.

Conclusions

The neoclassical utility function is usually viewed as only depending upon what an individual actually possesses and is not treated as a function of goals. But individuals clearly use goals in many of their decisions. Indeed evolutionary economics suggests that humans should have evolved to have utility functions that tend to support the goal of the survival of the human species. These considerations provide plausible motivations for Gill's proposed goal-dependent utility function.

This paper shows how the neoclassical utility function—as formulated by von Neumann and Morgenstern—can be explicitly written as the probability of achieving a goal. Since von Neumann and Morgenstern derive their result from axioms of rationality, this indicates that rationality leads to behavior that appears goal-oriented and to utility functions that are goal-based. In addition to suggesting ways of explaining paradoxes, this representation also highlights new connections between such unrelated fields as finance, medicine and complex systems analysis. We hope this paper encourages further work on Gill's profoundly important thesis.

References

Borch, K. (1968). Economic objectives and decision problems. *IEEE Transactions on Systems Science and Cybernetics, 4*(3).266-270.

Bordley, R., & LiCalzi, M. (2000). Decision analysis using targets instead of utility functions. *Decisions in Economics and Finance, 23*(1), 53-74.

Bordley, R., & Kirkwood, C. (2004). Multiattribute preference analysis with performance targets. *Operations Research, 52*(6), 823-835.

Bordley, R. (2009). The Hippocratic Oath, effect size, and utility theory. *Medical Decision Making, 29*(3), 377-379.

Castagnoli, E., & LiCalzi, M. (1996). Expected utility without utility. *Theory and Decision, 41*(3), 281-301.

Gill, T. G. (2008). A psychologically plausible goal-based utility function. *Informing Science: The International Journal of an Emerging Transdiscipline, 11*, 227-252. Retrieved from http://www.inform.nu/Articles/Vol11/ISJv11p227-252Gill220.pdf

Oden, G., & Lopes, L. (1999). The role of aspiration level in risky choice: A comparison of cumulative prospect theory and SP/A theory. *Journal of Mathematical Psychology, 43*(2), 286-313.

Schervish, M, Seidenfeld, T., & Kadane, J. (1990). State-dependent utilities. *Journal of the American Statistical Association, 85*, 840-847.

Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior.* Princeton, NJ: Princeton University Press.

Biography

Robert Bordley earned a PhD and MS in operations research and an MBA in finance from the University of California, Berkeley. He also holds an MS in systems science, a BS in physics, and a BA in public policy from Michigan State University. He is an INFORMS Fellow and winner of five major awards from General Motors Corporation as well as a Publication Award from the Decision Analysis Society. Bordley has published 75 papers in a variety of journals in marketing, operations research, statistics, and economics. Bordley has served as program chair for the Statistics in Marketing Section of the American Statistical Association and as former chair of the Risk Analysis Section. He has twice been vice president of the Production and Operations Management Society and a council member of the INFORMS decision analysis society.