A PSYCHOLOGICALLY PLAUSIBLE GOAL-BASED UTILITY FUNCTION

T. Grandon Gill, IS&DS Dept., CIS1040, University of South Florida, Tampa FL 33620
ggill@cobal.usf.edu, 813-974-6755

INTRODUCTION

The notion of a utility function, that is to say a function that captures human preferences and thereby explains human choices, is widely used in many business-related fields. In economics, it plays a foundational role in the mathematical derivation of general equilibrium theory. In the decision sciences, it is central to rational models involving risk and time. In management and finance, it serves as an important underpinning to a wide range of models, such as those employed in agency theory and prospect theory.

Despite its ubiquity, the shape and properties of the utility function remains a subject of considerable research (and debate). Even in its simplest form, a single variable and one time period, numerous departures from apparently rational behavior as predicted by neoclassical utility functions have been observed in experimental settings. Henceforth, we refer to theses as anomalies. As additional dimensions, such as time or non-independent products, are added to the functions, their mathematical tractability (and therefore usefulness in formal derivations) quickly declines. Where utility functions are used to explain individual behaviors in organizations, all manner of non-measureable quantities, such as individual "power", need to be incorporated as arguments. Thus, while the utility function may serve a highly useful purpose in a theoretical context, it is not clear that its practical benefits are nearly as significant.

The purpose of the present paper is to synthesize existing psychological research into an alternative formulation of the utility function that may serve to address some of the anomalies and inconsistencies common to many existing utility functions. The paper begins by summarizing existing models of utility and then reviews key findings of the goal-setting literature. It then formulates a utility model defined strictly in terms of individual goals. Finally, it considers how predictions of the goal-based model might differ from existing models.

UTILITY THEORY

Utility generally refers to an individual state that combines notions of satisfaction, usefulness and rationality. In economics, utility is often presented as a function of various goods and services $(x_1,x_2,\ldots,x_n)$ or commodities [3] consumed in a period along with resources saved $(s_1,s_2,\ldots,s_n)$, e.g.,

$$U(x_1,x_2,\ldots,x_n, s_1,s_2,\ldots,s_n)$$

Neoclassical economic theory is built upon the assumption that any rational economic agent can be modeled as a utility maximizer. Conceptually, we may view the utility maximization process as a choice process wherein a decision-maker starts with a set of saved resources supplemented by earnings, considers how those resources may change as a consequence of possible actions available to him or her, then selects that action leading to the highest utility outcome.

Because utility maximization is axiomatic to theories that postulate rational decision making, the underlying shape of individual utility functions has been researched extensively. Such research generally falls into one (or more) overlapping categories: the impact of risk on utility, the impact

Table 1: Anomalies in utility models of rational decision making

Anomaly	Description	Example Reference
Preference Anomalies		
Framing effects	Individuals often express strong preferences when making choices between alternatives that are, in fact, identical.	[5], p. 11.
Endowment effect	Upon acquiring an object, we value it far more than we did prior to acquiring it, leading to asymmetry between buying and selling prices.	[1], p. 130.
Anchoring	When establishing preferences, we typically anchor them to an available reference, even when that reference makes no sense (e.g., the last digits of a social security number).	[1], p. 28.
Preference reversal	When presented with two lotteries, individuals may prefer the less risky one yet be willing to pay more to participate in the riskier one.	[10]
Illusion of control	Individuals perceive that they exert far more control over random events than they actually do.	[15], p. 231.
Cash effect	Experimental subjects react differently to cash rewards than to rewards easily convertible to cash.	[1], p. 220.
Availability bias	Judgments of likelihood are unduly influenced by the individual's ability to recall specific examples, which may or may not be representative.	[27], p. 163.
Expected value insensitivity	When presented choices of certainty and near certainty, we tend to prefer certainty even where payoff differences are substantial. As probabilities get very low, we focus on size of the payoff rather than its expected value.	[12]
Temporal Anomalies		
Common difference	Identical payoff separated by a fixed time period may change preference depending upon when the payoff is started. For example, you may prefer $3000 now to $4000 in a year, but also prefer $4000 in year 4 to $3000 in year 3.	[19], p. 120.
Absolute magnitude effect	Ratios of preferences may change as payoffs change. For example, an individual may prefer $15 now to $60 in a year, but prefer $4000 in a year to $3000 now.	[19],, p. 121.
Gain-loss asymmetry	Losses are be discounted at a greater rate than gains. For example, a study found an individual be indifferent to a gain of $10 now to $21 in a year, but be indifferent to a loss of $10 now and a loss of $15 in a year.	[19], p. 122.
Delay-speedup asymmetry	Significant variations were found with respect to an individual's willingness to speed up consumption and delay consumption. In other words, they had to be paid more to move from t1 to t2 than they were willing to pay to move from t2 to t1.	[19], p. 124.
Variance from future expectations	Individuals tend to consistently overestimate actual utility gains and losses when considering future events.	[8], p. 102.
Sequence effect	Individuals generally exhibit a preference for a sequence of increasing payoffs that total to the same amount as decreasing payoffs.	[23], p. 435

of time on utility and the behavior of multi-attribute utility functions. Research into risk has focused on better understanding how people trade off costs and payoffs that include an element of randomness, either quantified (risk) or unspecified (uncertainty or ambiguity). One surprising result of this research has been the discovery of numerous anomalies that appear to call into question the rationality of actual human behavior (see Table 1). A similar stream of research has identified anomalies in utility tradeoffs over time.

Utility preference research investigating the impact of both risk and time generally focuses on a single attribute, money, as the underlying source of utility. Another stream of research, however, focuses on multi-attribute utility functions. The challenge presented by this research is that of mathematical tractability. Generally, multi-attribute functions pose little problem provided that the utility from different attributes is linearly additive, i.e.,

$$(2) \qquad U(x_1,x_2,\ldots,x_n) \equiv U_1(x_1)+ U_2(x_2)+\ldots+ U_n(x_n)$$

Where interdependencies exist, however, the mathematics of utility maximization becomes much more complicated. For example, if an individual's utility from acquiring a bottle of foreign beer simultaneously depends upon the ability to acquire a bottle opener, then the utility of neither can be computed independently (particularly if no such opener is required for a domestic beer). One way of handling such problems mathematically is to include cross products in the utility function [13]. Unfortunately, as the number of attributes and interdependencies in the utility function grows, any closed form solution using this approach becomes impractical [4].

MOTIVATION AND GOAL SETTING THEORY

When viewed in the context of the management field, expected utility is essentially a measure of a state's motivational potential. Stated another way, the underlying axiom of utility theory is that an individual will be motivated to make the decision (or set of decisions) yielding the state of maximum expected utility. Motivation, however, has been the subject of many investigations within the management literature—most of which bear little or no surface resemblance to neoclassical utility theory. Perhaps the most important of these research streams is goal setting theory, some findings of which are summarized in Table 2.

Table 2: Important findings in goal-setting theory

Finding	Description	Example References
Priming	By unconsciously stimulating certain aspects of a goal, for example through having subjects read through a list of positive and negative words prior to participating in an experiment, substantially different goal performance and satisfaction may be realized.	[26], p. 1175
Approach vs. Avoidance	Eagerness increases as you near completion of an approach goal, vigilance increases as you near completion of an avoidance goal.	[7], p. 1129.
Specificity	Goals that are specific and difficult lead to better performance than a vague goal or no goal at all.	[16], p. 332.
Mastery vs. Performance	Goal striving leads directly to enjoyment for mastery (learning) goals, and indirectly—through mental focus—for performance goals.	[18], p. 262
Commitment	Performance grows with goal commitment, especially for difficult tasks.	[14], p. 886
Proximity	People exert more effort to achieve goals where progress has already been made, referred to as the "endowed progress effect".	[21], p. 510
Difficulty	Increased goal difficulty leads to increased motivation and performance for approach goals, provided the goal is achievable, but less so for avoidance goals.	[11], p. 377
Participation	Goals provide an important motivational regardless of whether or not the individual has participated in setting the goals. In many cases, participation doesn't appear to matter.	[17], p. 416
Goals vs. Needs	Goals are better predictors of performance than personal achievement needs.	[2], p. 264

A GOAL-BASED UTILITY FUNCTION

As described earlier, motivation can be characterized as the impetus to maximize expected utility. If one accepts the findings of goal setting research it therefore stands to reason that a utility function could be constructed using goals, rather than other attributes (such as money), as its arguments. The validity of such a function would depend on several axiomatic assumptions, namely:

1. At any given time, an individual holds a set of conscious goals that can be articulated. Such goals may derive from many sources, including individual needs and external sources, such as the well-being of friends and family, the goals of organizations and communities with which the individual is affiliated, and even broader sources, such as spiritual goals, national goals and environmental goals.
2. Utility derives from two processes: the direct satisfaction of needs and progress in the pursuit of the individual's goals, both in absolute terms and relative to expectations. For the purpose of the model, needs are treated as concrete biological goals [22](p. 30), allowing all utility to be treated as arising from goal-based sources.
3. Attending to all goals concurrently would exceed the individual's processing capacity by many orders of magnitude. As a consequence, mechanisms for prioritizing, selecting, scheduling and refining active goals are critical elements of reasoned behavior. These come in two forms: unconscious processes that establish preferences and heuristics—both specialized and general-purpose—that tend to be consciously applied. We refer to the latter as meta-reasoning behaviors.
4. As a result of axiom (2), we may assume that the decision to engage in explicit meta-reasoning behaviors is also goal driven. Goals of this type will be referred to as meta-goals.

The actual goals driving the goal-based model would result from a number of sources. Goals are created as responses to underlying needs and drives, consistent with the goal-setting model. Lower level need-based goals—such as those towards the bottom of Maslow's hierarchy—would routinely arise from fundamental biological drives (e.g., hunger, sex, security, social interaction) as a consequence of planning for future satisfaction of those needs. Where planning is not required, a direct path from need satisfaction to utility can occur. The model would predict, however, that the individual's utility gain from directly satisfying a need versus progressing towards a conscious goal derived from that same need could be quite different. This is consistent with empirical findings that reasoning about a goal changes its utility [6].

Higher levels of the need hierarchy, much more uniquely human, would drive the formation of achievement goals, such as career advancement. These would tend to be conscious in form. Meta-reasoning processes, invoked by an evolved drive to reduce uncertainty, would focus on goal identification and planning-related activities such as prioritization and scheduling. For example, the environment may present the individual with an unexpected opportunity stimulus. In some cases, a special-purpose script—conscious or automatic, acquired through past experience or education—for translating that stimulus into a goal and creating an action plan may exist. Because such an opportunity can come in many shapes and forms, however, each decision maker must also possess a set of general purpose heuristics for recasting novel stimuli into utility preferences and action plans. As previously noted, general-purpose preference heuristics have been the subject of considerable study. Where preference determination is

required, unconscious and conscious processes may both be involved. In fact, it has even been shown that for certain types of choice problems relying on unconscious processes may lead to better results than applying available heuristics [6][9][28]. Planning heuristics, on the other hand, would nearly always be conscious. These would include general purpose strategies such as decomposing a complex goal into subgoals [24] and various search heuristics [20].

A goal-based utility function differs from neoclassical utility functions—which are not precluded from incorporating goals as attributes [25]—in that it proposes that an understanding of individual goals and meta-goals is a prerequisite to meaningful explanations and predictions of preferences and decisions. In the goal-based utility model, utility would be a function of at least three distinct elements:

1. The attractiveness of the goal outcome (roughly equivalent to the valence in expectancy theory), which in turn would be influenced by goal type (e.g., learning vs. performance, approach vs. avoidance goals, specific vs. general goals will tend to have different profiles).
2. The perception that he or she is making progress towards the goal (including activities that create a more concrete path towards the goal, such as planning). Where the individuals own actions are seen as being instrumental in achieving such progress, the utility would be even greater, as the desire to exercise control is another example of a fundamental human need.
3. The degree of commitment to the goal, as approximated by its goal level (Figure 1), with goals moving up (towards attended) increasing utility, and moving down (towards impossible) reducing utility.

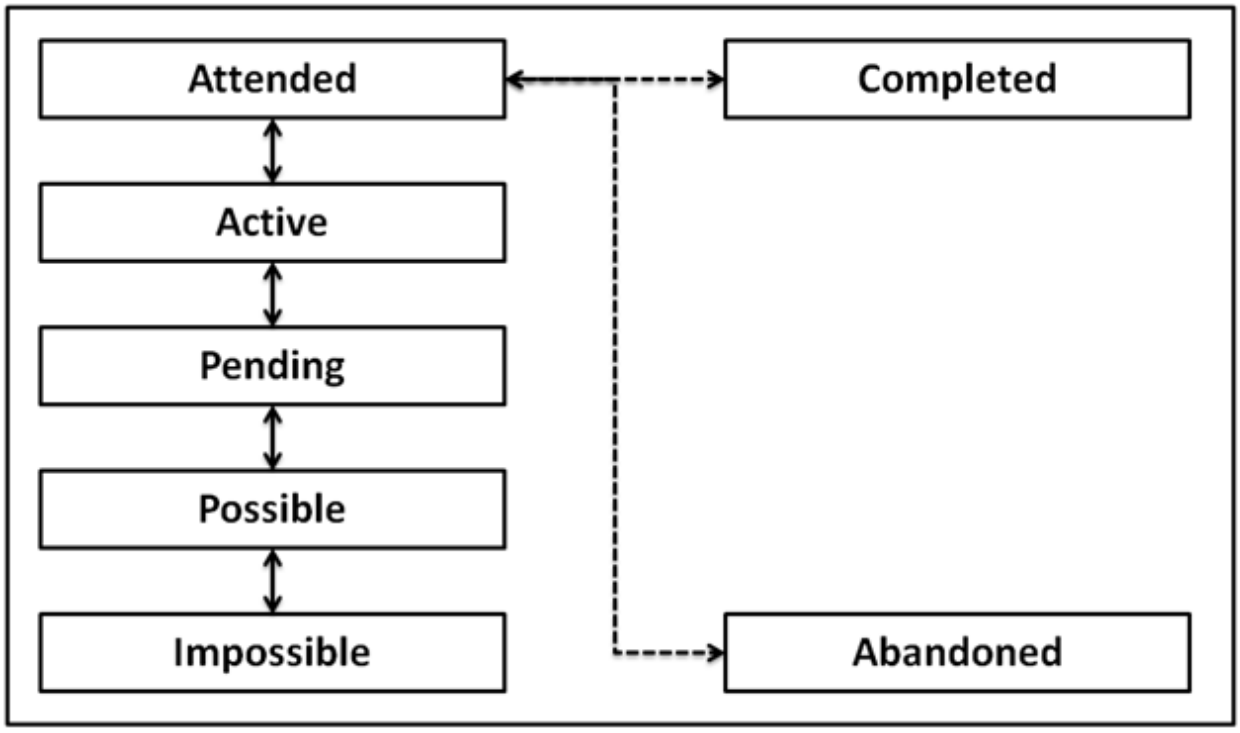

Figure 1: Goal Levels

CONCLUSIONS

Perhaps the greatest difference between the neoclassical and goal-based utility models can be characterized in terms of opacity of decision processes. The neoclassical model is constructed very much according to the behaviorist philosophy in which a stimulus is provided in the form of a set of alternatives, response is determined in the form of a choice, and relatively little attention paid to the decision maker's self-reported internal thought processes. The advantage of such an approach, particularly for macro-economic purposes, is that errors in predicting individual decisions are likely to get aggregated away. Weaknesses in neoclassical models only become embarrassingly apparent when specific individual preferences and behaviors are the object of prediction.

The goal-based model, in contrast, presumes continual adoption, modification, completion and abandonment of specific goals. It therefore complements the neoclassical model by being entirely decision-maker directed. It involves analyzing information processing activities within the decision maker's head, making it a cognitive (rather than behavioral) model. Indeed, without knowing the decision maker's individual goals, it would be nearly impotent in its ability to make any useful predictions. Its advantage would, in principal, therefore lie in interpreting behaviors and making predictions at the individual level.

The significance of better understanding utility has implications that extend beyond the immediately obvious. As an example, incentive-based compensation schemes are often presented as an approach to better align executive goals with those of shareholders and away from inappropriate self-interested goals. The goal-based model would predict that achieving such an alignment *will* lead to greater attendance of the desired financial performance goals, along with personal financial goals. By the rule of selective attendance, however, such executives would also then be expected to experience lesser utility from progressing towards other goals—which, unfortunately, are likely to include such important non-economic motivations such as maintaining personal integrity, ensuring accounting accuracy or achieving fairness. Furthermore, this predicted effect would tend to be largely independent of the amount of compensation involved, above a certain threshold required to drive initial attendance towards the financial goals. Thus, as a paradoxical consequence of priming executives to focus exclusively on maximizing shareholder wealth, we might well be opening the door to a whole class of unscrupulous behaviors—many of the very behaviors that we had hoped to deter through establishing the objectively-based compensation scheme. As Ariely aptly states: "just thinking about money makes us behave the way most economists think we behave" [1](p. 75).

Naturally, considerably more confirmatory research needs to be conducted before we begin recommending the re-engineering business systems according to the predictions of the goal-based utility model. In the final analysis, the purpose of this paper has been to leave the reader with three key points. First, the major body of consistent research on the effects of goal-setting virtually demands that our conception of utility incorporates the satisfaction we experience when progressing towards goals. Second, that the proposed model is specified in such a way that many testable hypotheses can be generated to confirm or refute it. Third, the topic is of sufficient practical importance that conducting such confirmatory research should be a priority.

References are available from the author.