Utility Analysis for Decisions in Human Resource Management

Working Paper 88-21

John W. Boudreau
Associate Professor

NYSSILR--Cornell University
393 Ives Hall
Ithaca, New York 14851-0952
(607) 255-2273

Draft of a Chapter for the
Handbook of Industrial-Organizational Psychology
Edited by Marvin Dunnette
to be published by Consulting Psychologist Press
Palo Alto, California

December, 1988

This research was carried out with support from the U.S. Army Research Institute, contract SFRC #MDA903-87-K-0001. the views, opinions and/or findings contained in this chapter are those of the author and should not be construed as an Official Department of the Army policy, or decision.

jrg

Introduction

The questions studied by Industrial/Organizational (I/O) psychologists are closely linked to the decisions facing managers of people in organizations. Whether they be line managers, human resource management staff, or organizational psychologists, managers of human resources must make decisions about issues affecting the employment relationship--hiring, training, compensation, performance appraisal, and so on--that draw on theories of human work behavior. Analogously, I/O psychologists, as well as other social scientists, find the organizational environment a rich source of information for advancing knowledge and testing employment-related theories. Both scientists and managers benefit from the knowledge gained about the behaviors of individuals in the work place, who can then search for ways to apply that knowledge to achieve individual and organizational outcomes of efficiency and equitable employment.

The similarity of interests between I/O psychologists and human resource management (HRM) professionals has produced some close collaborative relationships, e.g., the many psychologists who consult for industry, conduct studies designed to support HRM decisions, or, through their work, influence the direction of employment policies. Still, the HRM functions of organizations typically lack the influence and visibility of other management functions such as marketing, finance, and operations. The literature for HRM professionals routinely laments the slow implementation of HRM programs in organizations, even though these programs have gained wide acceptance by scientists (cf. Jain & Murray, 1984), and they admonish and instruct these professionals to "sell" their programs by emphasizing their effects on attainment of organizational goals (Bolda, 1985, Fitz-Ens, 1984; Gow, 1985, Jain & Murray, 1984, Sheppeck & Cohen, 1985). With increased competition and evidence from the United States and abroad that competitive organizations are likely to manage their people differently, HRM personnel are more frequently expected to justify their contributions to the employer and to account for their existence.

One must question whether the lack of influence and slow implementation of HRM programs is a rational response by organizations. Could it be that behavioral theories and findings are relevant only to the scientific community and have such little relevance to organizational decisions and outcomes that they can be ignored by successful organizations? If the theories and findings are relevant, then how should they be communicated to decision makers? Do decisions that consider social science evidence produce greater organizational success, and, if so, are the successes great enough to justify the resources necessary to generate and apply the evidence?

This chapter will discuss utility analysis (UA), which attempts to answer such questions by focusing on decisions about human resources. Utility analysis refers to the process that describes, predicts and/or explains what determines the usefulness or desirability of decision options, and examines how that information affects decisions. In HRM and I/O psychology, the focus lies on decisions involving employment relationships and employee behaviors. Thus, I/O psychologists use the term utility analysis to refer to a specific set of models that reflect the consequences, usually performance-related, of programs designed to enhance the value of the work force to the employing organization.

Utility analysis offers great potential for enhancing the link between the theories and findings of I/O psychological research and the human resource decisions of organizational managers. To achieve this potential, however, UA research and applications must proceed from a framework that recognizes the broad effects of such decisions on the work force and the organization. Such a framework requires an expansive view of the decision tasks facing managers of people in organizations, a view that recognizes the contributions, limits and implicit assumptions not only of psychological models, but of models from other social sciences as well. The UA framework provides both a rationale and a significant new direction for an integration between the science and practice of I/O psychology and other scientific disciplines relevant to organizational employment decisions. This chapter is intended as a step toward such an integrative framework. Thus, it will not only review and describe UA theories and applications, but will propose new and integrative directions that have received little attention. UA research must certainly acknowledge the considerations of related disciplines such as economics, management and sociology. But as a true theory of organizational decision making, it provides a mechanism to go beyond simple acknowledgement, to achieve a mechanism for truly interdisciplinary approaches to employment decisions.

Chapter Outline

This chapter comprises ten sections. The first section introduces and establishes some fundamental concepts, including the nature of utility models, decision options, attributes and payoff functions. It shows where UA models fit within the broader domain of decision models. It further establishes some ground rules guiding subsequent sections.

The second section outlines the historical development of concepts integral to utility analysis, the roots of which can be traced to the earliest stages of I/O psychological research. Not only does this historical outline provide some basic concepts for those not familiar with UA research, it also identifies certain fundamental concepts and assumptions essential to understanding utility analysis, which are sometimes ignored or forgotten in more recent theoretical developments.

The third section summarizes findings from previous studies revealing the effect of I/O psychological interventions on work force consequences. The fourth section critically reviews the research topic commanding the greatest attention to date--measuring the dollar value of performance variability.

The fifth section examines UA research from the perspective of information theory, by examining the role of risk and uncertainty in decision making. Such a perspective suggests that UA models can improve decisions even when information is severely lacking. Methods for identifying risk and uncertainty are described, as well as a technique for identifying when additional information is valuable. The role of UA research in defining statistical and substantive significance is also discussed.

The sixth section presents enhancements to the traditional selection utility models. These include incorporating financial/economic considerations, "intangible" factors such as equal employment opportunity and affirmative action, and the role of "constituencies" (Tsui, 1984; 1987; Tsui & Gomez-Mejia, 1988) in evaluating the usefulness of HRM programs. This section also shows how UA research can link I/O psychology and labor economics. It suggests that UA offers a mechanism for truly interdisciplinary approaches to employment issues, but that this demands that UA models reflect economic considerations, stocks and flows.

The eighth section discusses the role of utility analysis in describing consequences of programs that affect the "stock" of existing employees by altering the characteristics of the work force or work situation. Recent research is reviewed, suggesting implications for extending utility analysis research to important new areas.

The ninth section presents a unified utility model reflecting outcomes of HRM decisions affect the composition of the work force by changing the "flows" of employees into, through, and out of organizations--an employee movement utility model. Important links are proposed between recruitment, selection, turnover, and internal staffing. Empirical simulation analyses are described that suggest that the actual consequences of HRM decisions are likely to reach far beyond those reflected in current models addressing only the consequences of selection. It demonstrates the need for a fully integrated framework for considering the consequences of changing both the stocks and flows of employees, which can lead to greater synergy in planning and implementing employment programs.

Finally, the tenth section presents a matrix to guide future UA research, emphasizing the need to move beyond selection models and measurement issues, and toward a broader understanding of HRM program decision making.

Concepts and Definitions

Utility Analysis as a Subclass of Multiattribute Utility Analysis

Multiattribute utility (MAU) models are "decision aids" (edwards, 1977; Einhorn & McCoach, 1977; Einhorn, Kleinmuntz & Kleinmuntz, 1977; Fischer, 1976; Huber, 1980; Keeney & Raiffa, 1976) that provide tools for describing, predicting and explaining decisions. MAU models share certain characteristics and requirements. To apply such models, one must:

(1) Identify a set of *decision options* that represent the alternative programs or courses of action under consideration;

(2) Identify a set of *attributes* that reflect the characteristics of the options that are important because they represent the things that matter to the decision makers and/or the relevant constituents.

(3) Measure the level of each attribute produced by each option using a *utility scale* for each attribute;

(4) Combine the attribute values for each option using a *payoff function* reflecting the weight given each attribute and combination rules for deriving an overall total utility value for each option.

Insert Table 1 Here

Table 1 illustrates an extremely simple application of MAU analysis. Suppose productivity is below desired levels among sales people. Two *decision options* might be identified, involving two different training programs called Program A and Program B. Three *attributes* are of interest: (a) Effects on sales levels, (b) Resources required to develop and implement the program; and (c) Effects on sales person job satisfaction. Attributes (a) and (b) use a *utility scale* of dollars, while Attribute (c) uses a rating scale from 1 to 7. The *payoff function* consists of multiplying the level of attributes (a) by 1, multiplying the level of attribute (b) by -1, multiplying the level of attribute (c) by 3,000, and adding the results to produce a total utility value. We could construct a Multiattribute utility matrix like that shown in Table

1, with the cells of the matrix containing the expected level of each attribute for each option, and the total utility values below each option computed using the payoff function. Although Program B has the higher first-year dollar payoff, the high weight given to attribute (c), Job Satisfaction, combined with Program B's lower Job Satisfaction cause it to attain a lower utility value than Program A, and thus to be less preferred. Obviously, MAU models can encompass a variety of decision options, numerous and diverse sets of attributes reflecting many different constituents, and very complex payoff functions, but they generally share the characteristics shown in the simple example of Table 1.

MAU models can assist decision makers in overcoming "limits on rationality" (March & Simon, 1958) by providing a simplified, structured framework within which to consider a number of decision options. Huber (1980, pp. 61-62) identifies five advantages of MAU models over less systematic and structured decision systems:

(1) Because they make explicit a view of the decision situation, they help to identify the inadequacies of the corresponding implicit, mental model;

(2) The attributes contained in such models serve as reminders of the information needed for consideration of each alternative;

(3) The informational displays and models used in the mathematical model serve to organize external memories;

(4) They allow the aggregation of large amounts of information in a prescribed and systematic manner; and

(5) They facilitate communication and support to be gained from constituencies.

As a subclass of MAU models, UA models also serve as decision aids, and can provide the advantages listed above. Unfortunately, very little theoretical or empirical research has approached utility analysis from this decision-making perspective. Nonetheless, a keen appreciation of the role of UA models in the decision process suggests some very different research questions and directions. These will be emphasized throughout the chapter. Unlike the generic MAU model describe in Table 1, UA models focus on a particular type of decision option, a restricted set of attributes, and a defined mathematical formula for attribute weights and combination rules. The next sections examine these MAU components, and how they apply to UA models.

The Decision Options: HRM Productivity-Enhancement Programs

Any MAU model requires a focus of analysis--the decision options considered. For example, an MAU model for deciding where to build a new hospital might focus on options reflecting different types of facilities, combined with different locations, combined with different service offerings. Each combination would constitute a decision option. Utility analysis has focused on HRM programs designed to enhance work force productivity. Such programs include selection testing, recruitment, training, and compensation--all of which affect the organizational value of the work force, whether they are explicitly chosen using decision models or evolve implicitly over time (Milkovich & Boudreau, 1988). Utility analysis involves describing, predicting and explaining the consequences of such program options, their desirability, and the decision processes leading to choices among them. Thus, while the focus of UA is more specific than generic MAU models, it covers a wide array of options relevant to organizational goals. As we shall see, the majority of UA research has focused only on selection programs, but we now have the theoretical models to apply to virtually any HRM program.

Decisions about individuals versus decisions about programs. Utility analysis models might seem to focus on decisions about individuals, rather than programs. For example Cascio (1980, p. 128) stated, "all personnel decisions can be characterized identically. In the first place there is an individual about whom a decision is required. Based on certain information about the individual (for example performance appraisals, assessment center ratings, a disciplinary report), decision makers may elect to pursue various alternative courses of action." In MAU terms, the decision options are different courses of action for each individual.

However, closer examination shows that UA models are intended to apply to decisions about the *programs* that guide the countless decisions about individuals made by human resource managers. The options under consideration are the procedures, rules or "strategies" (Cronbach & Gleser, 1965, p. 9) meant to be used with many individuals, and evaluated by their "total contribution when applied to a large number of decisions" (p. 23). Decisions about whom to hire depend on what programs of recruitment and testing have been chosen to generate applicants and information about them. Decisions about how much to pay individuals depend on what compensation programs and rules have been chosen for that work force. Decisions about assigning individuals to new jobs depend on what career development and training programs have been chosen to generate skills and forecast future needs. Thus, UA models focus on the more strategic and tactical decisions about programs, rather than the operational decisions about each individual.

Because program decisions affect many individuals throughout their tenure with the organization, the impact of even a single program decision on future work force consequences can be quite large. A selection program that affects the hiring decisions for 1,000 people, each of whom stays for 5 years affects 5,000 person-years of organizational behavior. If a more correct program decision produces even a modest work force quality increase of $10 per person-year, it's impact can be $50,000. Of course, this also suggests that the consequences of wrong decisions have large potential negative effects. Utility analysis uses information from social and behavioral sciences to attempt to improve such important decisions.

Two types of programs addressed by UA models. It is useful to group the variety of HRM programs that can be addressed by UA models according to whether they affect employee "flows" or employee "stocks". First, programs affecting employee *flows* change the composition or membership of the work force through "employee movement" (Boudreau, 1988; Boudreau & Berger, 1985a, 1985b; Milkovich & Boudreau, 1988). For example, *selection programs* allow additions to be made to the work force, *retention programs* determine which employees are retained when separations take place; and *internal staffing programs* determine which employees move between positions within an organization (Milkovich & Boudreau, 1988, Chapters 10-13). UA models applied to such programs focus on the process used to determine which individuals are chosen to move or to remain, and the program's consequences reflect the effects of having a different set of employees in the work force. UA models are typically applied to decisions about this type of program, with external selection programs receiving the greatest attention.

Second, programs affecting the employee *stock* change the characteristics of the existing set of employees, in their current positions. For example, *training programs* operate by altering knowledge, skills, attitudes, or other employee characteristics; *Compensation and reward programs* operate by altering the relationship between behaviors/outcomes and rewards; *Performance feedback and goal setting programs* operate by altering employee perceptions of the consequences of their behaviors. Such programs work to the extent that they lead to different behaviors by existing employees, which lead to more valuable organizational outcomes. UA models address decisions about such programs by focusing on options representing different kinds of programs affecting the stock of existing employees.

The Attributes of Programs In UA Models

Once a set of decision options is defined, MAU models specify the set of attributes reflecting the outcomes of concern to the decision makers and relevant constituents, and the level of each attribute achieved by each decision option. For example, the decision about where to build a hospital might include attributes as diverse as the environmental impact of the facility, speed of treatment in emergencies, and impact on local property values, reflecting the concerns of constituents as diverse as community planners, potential patients, nearby property owners and the future medical staff.

UA models focus on decisions about HRM programs, so the attribute set is more focused, but still quite broad. Cronbach & Gleser (1965, p. 22) defined the attribute domain as "all the consequences of a given decision that concern the person making the decision (or the institution he represents)." HRM program attributes may be placed in two categories--efficiency and equity (Milkovich & Boudreau, 1988): *Efficiency* attributes reflect the organization's ability to "maximize outputs while minimizing inputs", such as labor costs, job performance, sales volumes, revenues, profits, market share and various financial/economic indicators of organizational strength. *Effectiveness* attributes reflect the "perceived fairness" of organizational procedures and outcomes, such as employee attitudes, labor relations, minority and female representation, compliance with legal requirements, and community relations.

To date, most UA research and applications have focused on a very small set of efficiency-related attributes reflecting the productivity consequences of HRM program decisions. Although UA models can become mathematically complex, all existing UA models reflect just three basic attributes (Boudreau, 1984c, 1986, 1988; Boudreau & Berger, 1985a, 1985b; Milkovich & Boudreau, 1988):

(1) *Quantity*; the number of employees and time periods affected by the consequences of program options;

(2) *Quality*; the average effect of the program options on work force value, on a per-person, per-time-period basis;

(3) *Cost*; the resources required to implement and maintain the program option.

The program options addressed by UA models encompass a potentially large set of attributes reflecting both efficiency and equity, but existing model development has focused on a subset of the efficiency-related attributes reflecting program costs and employee productivity. Thus, like all models, UA models simplify reality by omitting or ignoring some factors. Models, by definition, are deficient because it is impossible to accurately reflect all the potential attributes affected by decisions. As we shall see, examining the nature of the attributes that are and should be included in utility models is one of the most critical issues facing UA research. Defining the domain of appropriate attributes offers fruitful opportunities for further debate and development. We will discuss these opportunities in some detail as we review existing research.

The Utility Scale for Attributes in UA Models

With the attributes identified, an MAU model must assign a value for each attribute in each decision option. This requires establishing a utility scale for each attribute, as well as determining the particular level of each attribute associated with each decision option. For example, in deciding where to locate a new hospital, the attributes are quite diverse (e.g., environmental impact, speed of treatment, facility cost, community satisfaction, etc.), and might be measured in units as diverse as dollars, time, number of complaints, ratings or rankings.

UA models focus on HRM programs, and therefore face a more limited set of attributes. Yet, even the relatively simple example in Table 1 had attributes measured in dollars (Costs and Productivity) as well as ratings (Job Satisfaction). UA models can potentially include a variety of efficiency and equity-related attributes, requiring diverse payoff scales. However, most UA models have focused primarily on productivity-related outcomes, striving to measure them in units relevant to managerial decisions. Attributes reflecting Quantity are usually measured in person-years, and those reflecting Cost are usually measured in dollars. The appropriate scale for the Quality attributes has been subject to some debate, as we shall see, but the majority of research has been devoted to scaling Quality in dollars per person-year.

Attaching a level of each attribute to each option often reflects a process using both subjective and objective information. When evaluating past programs, it may be possible to determine the actual levels of each attribute achieved by different options. But UA models are planning tools, used to anticipate future consequences and support current decisions, so attaching attribute levels involves predictions and forecasts. Indeed, one major motivation for UA models was to better express statistical forecasts in terms understandable to managers. The predictive nature of attribute measurement means that utility estimates possess uncertainty and risk. While uncertainty and risk take prominence in general MAU research, UA research has largely ignored them. As we shall see, mechanisms exist to promote further research in this important area.

The choice of attribute utility scales and derivation of attribute levels is important, and has received too little attention in UA research. Throughout this chapter we will highlight controversies where additional debate and research attention can be fruitful.

Combining Attributes Using a Payoff Function for UA Models

The fourth component of an MAU model is the payoff function, which specifies how the attribute levels are to be combined into an overall utility value. Deciding where to locate a new hospital might produce very diverse attributes measured on very different scales (e.g., dollars, time, and ratings/rankings). Payoff functions for such decisions must specify both the weights attached to each attribute level, as well as the rules for combining the weighted attribute levels to produce an overall utility value. Such rules might range from a simple numerical weighting and addition of the weighted values, to more complex non-linear weighting schemes and quadratic combination rules.

Because UA models focus on decisions about HRM programs, their attributes and payoff scales are more limited, and the payoff functions are often simpler. Still, any payoff function must reflect both the importance of each attribute and its underlying scale. The example in Table 1 adopted a relatively simple combination rule that takes the difference between increased productivity and costs, and then adds the Job Satisfaction level multiplied by 3,000. Obviously, the choice of weights and combination rules can have large effects of resulting utility values, and should reflect the values of the decision makers and relevant constituencies.

UA research has usually focused on productivity-related outcomes, and thus has adopted payoff functions reflecting dollar-valued productivity and program costs. The payoff function may be considered a variant of the cost-volume-profit models used in other managerial decisions to invest resources. The utility of an HRM program option is derived by subtracting Cost from the product of Quantity times Quality, with the program exhibiting the largest positive difference being preferred.

It is typical to refer to UA models as cost-benefit analysis models, and to categorize attributes as either Costs or Benefits. Simply put, Costs represent attributes that reduce overall utility values, while Benefits represent attributes that increase overall utility values. Depending on the decision, a given attribute (e.g., reduced employee separations) may represent either a cost or a benefit. Rather than attempt a classification, this chapter will proceed from the more general position that costs and benefits are defined by the attributes, their utility scales, and the payoff functions used to combine them. It is appropriate to question whether such a payoff function is adequate or even appropriate to UA research, and we will explore this issue at length.

Summary

UA research is a subclass of more general MAU research, and the structure of MAU models provides a useful framework for organizing and understanding UA models. As we have seen, UA models reflect a set of decision options, attributes, utility scales and payoff functions, just as any MAU model does. UA models have historically focused on a particular set of options (usually selection programs), attributes (Quality, Quantity and Cost), utility scales (Dollars) and payoff functions (Quantity times Quality, minus Cost). Measuring the payoff in UA research has been characterized as the "Achilles' heel" of UA research (Cronbach & Gleser, 1965, p. 121). As we have seen, such measurement reflects three MAU components: The attributes included; the utility scale used to measure them and attach a value to each option; and the payoff function specifying the combination rules across attributes. These components reflect implicit and explicit assumptions about the appropriate decision makers, constituents and consequences to be considered. Throughout the chapter, we will use these MAU concepts to organize and analyze existing and needed future UA research.

We have also seen that UA models, like all models, strike a balance between simplicity and realism. All UA models are deficient by definition, and much research debate has centered on whether and how to reduce that deficiency. But we will never develop a UA model that completely reflects all relevant attributes with perfect accuracy. Does this mean that UA research is unlikely to provide any real information about the effects of HRM program decisions on organizations? If the ultimate objective of UA is to *measure the impact* of program decisions on organizations, then the answer might be "yes", and we could declare a moratorium on UA research. However, like all MAU models, UA models are decision aids, not just measurement tools. A decision aid's usefulness lies in *its ability to describe, predict, explain and improve decisions.* Such value is assessed by asking whether the model allows the best decision to be made with the given body of information, whether it helps to determine if gathering more information would permit better decisions, and whether it helps to determine how much different decision procedures contribute to decision quality (Cronbach & Gleser, p. 21). Depending on the cost and value of the next best alternative decision aid, even a very deficient or inaccurate UA model might prove effective in improving decision processes or outcomes. Thus, this chapter will approach UA research less from a measurement perspective and more from a decision making perspective.

Historical Development of Utility Analysis Models[1]

Though utility analysis is applicable to virtually every HRM program decision, present models resulted from a concern with selection (and later, placement or classification) decisions. Indeed UA models can be characterized as responses to the inadequacies of traditional measurement and test theory in expressing the usefulness of tests.

> "The traditional theory views the test as a measuring instrument intended to assign accurate numerical values to some quantitative attribute of the individual. It therefore stresses, as the prime value, precision of measurement and estimation. The roots of this theory lie in surveying and astronomy, where quantitative determinations are the chief aim. In pure science it is reasonable to regard the value of a measurement as proportional to its ability to reduce uncertainty about the true value of some quantity. The mean square error is a useful index of measuring power. There is little basis for contending that one error is more serious than another of equal magnitude when locating stars or determining melting points: measurement theory is unobjectionable when applied to such appropriate situations.
>
> "In practical testing, however, a quantitative estimate is not the real desideratum. A choice between two or more discrete treatments must be made. The tester is to allocate each person to the proper category, and accuracy of measurement is valuable only insofar as it aids in this qualitative decision. ... Measurement theory appears suitable without modification when the scale is considered in the abstract, without reference to any particular application. As soon as the scale is intended for use in a restricted context, that context influences our evaluation of the scale." Cronbach and Gleser (1965, pp. 135-136).

Therefore, the history of UA will be discussed from a decision-making perspective, focusing on the contributions and implications of UA developments for describing, predicting explaining and enhancing decision processes and outcomes. Because the vast majority of research has emphasized the selection utility model, this will be the focus on the discussion. In this model, the option set involves using a test versus random selection (or choosing between two selection tests), and the utility value reflects only the effects of selection on the first job to which one group of selectees are assigned. Later sections will describe more recent developments that extend utility analysis beyond selection.

---

[1]This section emphasizes developments that set the stage for more recent research and future research directions. Much of this material is drawn from Boudreau (1987, in press). Other historical summaries can be found in Cronbach and Gleser (1965, chapter 4), Hunter and Schmidt (1982), Cascio (1982, 1987).

Defining the Payoff Based on the Validity Coefficient

Description of the Model. The attribute of selection tests that has the longest history is the validity coefficient, or correlation between a predictor measure and some criterion measure of subsequent behavior, usually expressed as $r_{x,y}$. Classical measurement theory suggested this concept as a measure of the "goodness" of a test in predicting subsequent behavior. In addition to the validity coefficient itself, two translations are most commonly cited (e.g., Cronbach & Gleser, 1965, chapter 4; Hunter & Schmidt, 1982), both of them lead to the conclusion that only relatively large differences in the validity coefficient produce important differences in the value of a test. First, one can translate the validity coefficient into the *index of forecasting efficiency* (symbolized as E) using Equation 1 below.

$$E = 1 - (1-r^2_{x,y})^{1/2} \quad (1)$$

This index, emphasized by early statistical texts (e.g., Kelley, 1923; Hull, 1928), indicates the proportionate reduction in the standard error of criterion scores predicted by the test, compared to the standard error of criterion scores predicted using only the group mean. Second, the *coefficient of determination*, or the squared validity coefficient appeared as early as 1928 in Hull's text, and reflects the proportion of variance shared by the predictor and the criterion.

Obviously, very large increases in validity are required to substantially increase these indexes. As Cronbach and Gleser (1965, p. 31) noted, "the index of forecasting efficiency describes a test correlating .50 with the criterion as predicting only 13% better than chance; the coefficient of determination describes the same test as accounting for 25% of the variance in outcome." Yet, correlations as high as .50 may be rare. In short, using these indexes, it appeared that very great improvements in testing would be necessary to have any substantial effect on organizational outcomes.

Evaluation From a Decision-Theory Perspective. As MAU models of a test's usefulness for decisions, such formulas are deficient. Only one attribute of the selection system is considered--the accuracy of prediction, expressed as the shared variance between two normally-distributed variables. From a decision-making perspective, the usefulness of a selection system depends on its ability to provide information that will improve decisions, where decision improvements are measured in terms of valued decision outcomes. Therefore, this model omits selection system attributes such as the quality of the existing selection system, the effect of the proposed selection system information on actual decisions, and the impact of those effects on valued consequences.

The utility scale for attaching attribute values to each option is a statistic that measures squared deviations from a predicted linear function. Thus, both positive and negative prediction deviations from the linear function are equally undesirable. This implies that a decision maker would consider overpredicting a qualified candidate's future performance just as costly as underpredicting it. In fact, of course, the important deviations from predictions are the ones that result in selection errors (i.e., selecting a candidate who should not have been hired, and/or failing to select a candidate who should have been hired). These models adopt an implicit payoff function that assigns equal value (or loss) to inaccurate predictions at all points in the predictor-criterion space (Wesman, 1953). Because there is only one attribute, there is no payoff function for combining different attributes. The statistic serves as the sole relationship, the smaller the area of the distribution lying in the false-positive or false-negative region); (2) lower selection ratios produce more improved success ratios (because lower selection ratios mean more "choosy" selection decisions, and the predictor scores of selectees lie closer to the upper tail of the predictor distribution); (3) base rates closer to .50 produce more improved success ratios (because as one approaches a base rate of zero, none of the applicants can succeed, so selection has less value; as one approaches a base rate of 1.0, all applicants can succeed even without selection, so selection has less value).

Evaluation from a decision-making perspective. The Taylor-Russell model reflects three attributes, rather than only the validity coefficient, but it still provides a limited description of selection program utility. Like its predecessors, this model ignores both the number of employees affected and the number of time periods during which that effect will last. the model's measure of Quality (proportion successful) is also troublesome because it does not reflect the natural units of value such as sales, productivity or reduced errors. Finally, the model excludes attributes reflecting program costs (Cascio, 1980, 1987), but cost differences will occur, especially as the selection ratio is changed by screening more/fewer applicants.

Scaling the base rate as a dichotomous criterion (i.e., success/fail) will often lose information because the value of performance is not equal at all points above the satisfactory level, nor at all points below the unsatisfactory level (Cascio, 1982, p. 135; Hunter & Schmidt, 1982, p. 235, Cronbach & Gleser, 1965, pp. 123-124, 138). More typically, performance differences exist within the two groups, so a continuous criterion scale could be more appropriate. Cascio (1982, 1987) suggests it may be more appropriate for truly dichotomous criteria (e.g., turnover occurrences), or where output differences above the acceptable level do not change benefits (e.g., clerical or technician's tasks), or where such differences are unmeasurable (e.g., nursing, teaching, credit counseling). Combining the attributes by assuming bivariate normality and linearity implied in the payoff function may also be unrealistic in some selection situations.

Some have proposed that the choice of the criterion cutoff is "arbitrary" (Cascio, 1982, p. 133; Hunter & Schmidt, 1982, p. 235; Schmidt, Hunter, McKenzie & Muldrow, 1979) because it is set by management consensus or because objective information on which to base such a decision is rarely available, and that changing this "arbitrary" cutoff will change the base rate, and thus substantially alter the conclusions from the model. If indeed there is no objective method of setting the performance cutoff, then the Taylor-Russell utility model is inappropriate. However, the concept of a criterion cutoff is not arbitrary, nor does the Taylor-Russell model imply that arbitrary changes in that cutoff are to be regarded as legitimate methods of enhancing the success ratio. Rather, the criterion cutoff (and the base rate it implies) should be based on the relationship between the selection situation (i.e., the level of minimally-acceptable criterion levels) and the applicant population (i.e., the proportion of the population that would exceed that level if hired). This concept is essential to evaluating the effects of recruitment on staffing utility, and should not be abandoned by labelling it "arbitrary."

Variations on the Taylor-Russell model. The models discussed next add program costs to the model and/or redefine the attribute utility scales to include dollar-scaled consequences of different selection mistakes. Cascio (1980, p. 35) noted that Smith (1948) provides a method of adjusting the Taylor-Russell results to reflect pre-existing selection ratios and validities. Technically, if current-employee characteristics are used as inputs to the model, this assumes that current employees are similar what would hypothetically be obtained if the criterion itself were used to select employees). Moreover, he used the principles of linear regression to demonstrate the relationship between the correlation coefficient and increases in a criterion (measured on a continuous scale). Brogden's logic serves as the basic building block for virtually all subsequent UA research.

Assuming a linear relationship between criterion scores (y) and predictor scores (x), the best, linear unbiased estimate of the criterion score associated with a predictor score is:

$$E(y) = A + B(x) \quad (2)$$

The intercept (A) and the slope (B) of this line reflect the linear relationship between x and y as well as the units in which each of them was originally scaled. However, because predictor and criterion scales vary from study to study, it is difficult to compare these parameters or to use them in a general model. However, if we transform both the y and x variables into standardized (Z-score) units (i.e., Z_x and Z_y), we can write Equation 2 as follows:

$$Z_y = (r_{x,y})(Z_x) \quad (3)$$

Therefore, if we knew the average standardized predictor score of a selected group of applicants (i.e., $\bar{Z}_x$), our best prediction of the average standardized criterion score of the selected group (i.e., $\bar{Z}_y$) would be the product of the validity coefficient and the standardized predictor score, as shown in Equation 4.

$$\bar{Z}_y = (r_{x,y})(\bar{Z}_x) \quad (4)$$

The validity coefficient was well established. One way to estimate the average standardized test score of the selected group would be to actually observe the value after applying a selection device. However, Kelley (1923) suggested that if one assumes that the predictor scores are normally distributed and that one ranks applicants by test score and selects from the top down, then the average standardized predictor score is a function of the proportion of the applicant population falling above the predictor cutoff score (i.e., the selection ratio). However, if one assumes the predictor is normally distributed, then Equation 4 holds only if one also assumes normally distributed criterion scores as well.

Brogden (1949, Equation 6) and Cronbach and Gleser (1965, p. 309) make use of this approach to derive their models. If we symbolize the "ordinate of the normal distribution" corresponding to the standardized predictor cutoff score as lambda (i.e., λ), and the selection ratio corresponding to the standardized predictor cutoff as SR (it has also been symbolized by the greek letter ϕ), then, Equation 3 can be re-written:

$$\bar{Z}_y = (r_{x,y})(\lambda / SR) \quad (5)$$

The "ordinate of the normal distribution" is an important variable, multiplicatively related to the average standardized predictor score, and a statistically sophisticated concept. It is sufficient, however, to