IMPLEMENTING STOCHASTIC DOMINANCE WITH RESPECT TO A FUNCTION

Robert P. King and Lindon J. Robison

Introduction

The expected utility hypothesis is the basis for much of the large body of theory concerned with decision making under uncertainty. It is the source of a general decision rule--expected utility maximization--which permits the synthesis of information on decision maker preferences and expectations in a manner that is both analytically elegant and intuitively appealing. Despite its wide acceptance as a theoretical tool, however, the usefulness of the expected utility hypothesis in the solution of practical problems has been limited by several important operational difficulties.

One particularly serious problem associated with the implementation of the expected utility hypothesis is that empirically estimated utility functions often prove to be unreliable representations of decision maker preferences (Robison and King). Sources of error include shortcomings in the design of elicitation interviews, failure to consider more than a single performance criterion by which choices are evaluated, and respondents' own lack of precise knowledge about their preferences. Despite such problems, a utility function, once estimated, is usually treated as though it were an exact representation of preferences when it is used to order alternative choices, and any absolute difference in the expected utilities associated with two possible action choices is taken as a clear indication that one is preferred to the other. As a result, inaccuracies in an elicited utility function can cause the rejection of an action choice that is actually preferred by the decision maker.

Imprecision in the measurement of decision maker preferences can be recognized explicitly in a decision analysis by using an efficiency criterion rather than a single valued utility function to evaluate alternative choices. An efficiency criterion is a preference relationship which provides a partial ordering of feasible action choices for decision makers whose preferences conform to certain rather general specifications. As such, an efficiency criterion can be used to eliminate some feasible choices from consideration without requiring precise information about preferences. First and second degree stochastic dominance are among the simplest and most commonly used efficiency criteria. Both are fully consistent with the expected utility hypothesis, and both require only minimal information about decision maker preferences. Unfortunately, however, neither is a particularly discriminating, evaluative tool. In an application of second degree stochastic dominance by Anderson, for example, 20 of 48 randomly generated farm plans were in the efficienct set. While single valued utility functions often exclude too many choices from the efficient set, then, these commonly used efficiency criteria often fail to exclude enough choices.

Robert King is an Assistant Professor in the Department of Economics at Colorado State University. Lindon Robison is an Assistant Professor in the Department of Agricultural Economics at Michigan State University.

Computational problems have also limited the applicability of the expected utility hypothesis in decision situations which require the consideration of a large number of possible actions. Mathematical programming techniques are commonly used in the analysis of complex decision problems, but a number of serious difficulties are encountered when such techniques are employed in the analysis of decisions made under uncertainty. These difficulties impose rather severe restrictions on the representation of decision maker preferences, on the nature and complexity of probability distributions associated with alternative choices, and on the types of decisions that can be analyzed. Quadratic programming, for example, requires that outcome distributions be normal and that decision makers have utility functions of the negative exponential or quadratic form.[1] Linear programming alternatives to quadratic programming, such as the focus-loss (Boussard and Petit), game theoretic (McInerney), and MOTAD (Hazell) models, do not require that outcome distributions be normally distributed, but they employ choice criteria which are not fully consistent with the expected utility hypothesis. Finally, all of these models impose rather serious restrictions both on the complexity of the stochastic processes that can be modeled and on the kinds of management strategies that can be evaluated. With regard to the first of these problems, it is particularly difficult in a mathematical programming framework to represent stochastic processes in which random variables interact in a non-additive fashion, as is the case when both prices and yields are random in a farm planning situation. It is also difficult to deal satisfactorily with the impact of stochastic resource constraint levels in a mathematical programming model.[2] With regard to the types of decisions that can be analyzed, choices under uncertainty often take the form of flexible strategies which make forthcoming actions contingent upon future events that the decision maker can observe but not control (Massé). Such strategies are not easily evaluated within standard linear programming models; and, as a result, rather unrealistic inflexible action choices are often the only ones considered.[3]

In this paper we introduce two related analytical techniques that help to resolve some of the problems identified above. Both grew out of an effort to implement stochastic dominance with respect to a function (Meyer, 1977a), a recently developed efficiency criterion which can be used to provide a more complete preference ordering than can be achieved with first or second degree stochastic dominance. After a brief explanation of this criterion, the first of these new techniques will be discussed. It is a procedure for constructing interval measurements of decision maker preferences--measurements that can be made as precise or imprecise as a particular decision situation dictates, so

---

[1]When used parametrically, quadratic programming requires only that decision makers be everywhere risk averse.

[2]Paris has formulated a symmetric quadratic programming model which, at least in part, alleviates this problem.

[3]Stochastic programming (Cocks, Rae) can be used to analyze such problems if they are relatively small. For large multistage problems, however, the input-output matrix quickly expands to an unmanageable size under this procedure.

that problems associated with both single valued utility functions and commonly used efficiency criteria can be avoided. Such measurements are required for the application of stochastic dominance with respect to a function. The second technique to be discussed is a computational procedure for identifying preferred choices which combines random search, simulation, and evaluation by the criterion of stochastic dominance with respect to a function. This model is flexible and computationally efficient, and it is well suited for the analysis of a wide range of practical decision problems. Finally, the use of these techniques will be illustrated with a simple example concerned with crop mix and land rental decisions under price, yield, and weather uncertainty.

Stochastic Dominance with Respect to a Function

Stochastic dominance with respect to a function is an evaluative criterion which orders uncertain action choices for classes of decision makers whose absolute risk aversion functions, $r(x)$, lie everywhere between specified lower and upper bounds, $r_1(x)$ and $r_2(x)$.[1] In effect, it provides an ordering based on an interval measurement of preferences which is analogous in many respects to a statistical confidence interval. Unlike evaluative criteria based on single valued utility functions, stochastic dominance with respect to a function does not require that a decision maker's preferences be specified exactly. Unlike other commonly used efficiency criteria, such as first and second degree stochastic dominance, which hold only for quite broadly and inflexibly defined classes of decision makers, stochastic dominance with respect to a function imposes no restrictions on the width or shape of the absolute risk aversion interval within which the decision maker's own absolute risk aversion function is said to lie.

Under this criterion, two alternative choices are ordered by identifying a utility function which conforms to the restrictions placed on the decision maker's absolute risk aversion function and which minimizes the difference between the expected utilities associated with the two choices.[2] Should the minimum of this difference be positive, the strategy with the higher expected utility is clearly preferred to the other by all decision makers whose absolute risk aversion functions conform to the specified constraints. If the

---

[1]The absolute risk aversion function (Arrow; Pratt) is defined by the expression:

$$r(x) = -u''(x)/u'(x)$$

where $u'(x)$ and $u''(x)$ are the first and second derivatives of a von Neumann-Morgenstern utility function. While such a utility function is unique only to a positive linear transformation, the absolute risk aversion function represents preferences uniquely.

[2]Optimal control methods are used to identify this function. Details of the solution technique are presented in Meyer (1977a) and an example showing how the solution is implemented is given in King.

minimum difference between the two expected utilities is negative, the two choices under consideration cannot be ordered by unanimous preference.

An Interval Approach to the Measurement of Decision Maker Preferences

Stochastic dominance with respect to a function is a remarkably flexible evaluative criterion which has the potential for being a valuable tool in applied decision analyses. Before this potential can be fully realized, however, an operational procedure must be developed for the determination of lower and upper bounds on a decision maker's absolute risk aversion function. Such a procedure is introduced in this section. It provides a means by which information revealed through a series of choices between carefully selected distributions can be used to establish lower and upper bounds on an individual's absolute risk aversion function. The degree of precision with which preferences are measured--i.e., the size of the interval between these lower and upper bound functions--can be specified directly in accordance with the characteristics on the problem under consideration. At one extreme the interval can be of infinite width; at the other extreme it can converge to a single line.

The procedure is based on the fact that, under certain conditions, a choice between two distributions defined over a relatively narrow range of outcome levels divides absolute risk aversion space over that range into two regions: one consistent with the choice and the other inconsistent with it.[1] The level of absolute risk aversion at which this division is made depends solely on the two distributions--i.e., their properties define the two regions. The decision maker's preferences, as revealed by his ordering of the two distributions, however, determine into which of these two regions his level of absolute risk aversion is said to fall. By confronting the decision maker with a series of choices between carefully selected pairs of distributions, the region of absolute risk aversion space which is consistent with his revealed preferences is repeatedly divided. With each choice a portion of that region is shown to be inconsistent with the decision maker's preferences, and the interval measurement of absolute risk aversion is narrowed. The procedure continues until a desired level of accuracy is attained. Upper and lower limits for the level of absolute risk aversion are determined at several income levels, and these values are then used to estimate upper and lower limits for the absolute risk aversion function over the relevant range of income levels.

A simple example should help to illustrate how the procedure works. Consider the three outcome distributions given in Figure 1. Each contains six possible outcomes which are said to have equal probability of occurring. It can be shown that distribution 1 is preferred to distribution 2 by all decision makers whose level of absolute risk aversion is greater than .0005 over the range of outcome levels covered by these two distributions. Distribution 2, on the other hand, can be shown to be preferred by all decision makers whose level of absolute risk aversion is less than .0001. The two distributions cannot be ordered by unanimous preference over the interval

---

[1]Concepts developed by Meyer (1977b) are used by King to demonstrate the validity of this statement.

1. Compare distributions 1 and 2 and indicate which one you prefer. If you prefer distribution 1, go to question 3; otherwise, to to question 2. _____

2. Compare distributions 1 and 3, and indicate which one you prefer. _____

3. Compare distributions 2 and 3, and indicate which one you prefer. _____

Distributions

1	2	3
2100	1000	1750
2400	2050	1950
2550	2650	2500
3100	3800	2750
3250	3900	3950
3450	5200	4000

Figure 1. A Sample Questionnaire for Interval Preference Measurement

(.0001, .0005), which can be termed a boundary interval in risk aversion space.[1] If a decision maker prefers distribution 1 to distribution 2 and if it is reasonable to assume that his absolute risk aversion function can be adequately approximated by a constant value over the range of outcome levels covered by these distributions, then it can be concluded that his level of absolute risk aversion over that range is not less than .0001, since there is unanimous preference for distribution 2 by decision makers less risk averse than .0001. Similarly, if he prefers distribution 2, it can be concluded that his level of absolute risk aversion is not greater than .0005. Preference for any one of the two distributions, then, identifies a particular portion of risk aversion space within which his own risk aversion function does not lie.

Boundary intervals can also be identified for distributions 1 and 3 and distributions 2 and 3. For distributions 1 and 3 the interval is (-.0001, .0001), with distribution 3 preferred below the boundary interval and distribution 1 preferred above it. For distributions 2 and 3 the interval is (.0005, .0010), with distribution 2 preferred below and distribution 3 preferred above.

Using this information as a guide, the series of questions at the top of Figure 1 was specified. They take a form similar to that of a programmed learning text. The decision maker is always asked to answer the first question, but which of the second two questions he answers will depend on the choice he makes in the first. Consider the case where the decision maker prefers distribution 2 to distribution 1 in responding to the first question. This implies that his level of absolute risk aversion is less than .0005. He is then directed to indicate his preference between distributions 1 and 3. If he prefers distribution 1, his level of absolute risk aversion is shown to be greater than -.0001. This combined with the information from the first question indicates that his level of absolute risk aversion lies on the interval (-.0001, .0005). Had he preferred distribution 3, his level of absolute risk aversion would have been shown to be less than .0001, which, when combined with the information from the first questions indicates that it lies on the interval ($-\infty$, .0001). Note that, given his response to question 1, the comparison required in question 3 would not have provided any new information. It could serve, however, as a consistency check, since preference for distribution 3 in this case would not be consistent with preference for distribution 1 in the first question.

Upper and lower bound absolute risk aversion functions constructed using this procedure for two decision makers are shown in Figure 2. Each is based on interval measurements made over four income ranges. Note that the slopes of the absolute risk aversion functions are not restricted. For decision maker A, the bounded interval slopes downward as income levels increase; while for decision maker B, it slopes upward and then downward. It should also be noted that the interval measurements for both decision makers contain negative as well as positive values at some income levels. When absolute risk aversion functions are derived from empirically estimated

---

[1]Second degree stochastic dominance with respect to a function (Meyer, 1977b) is used to identify boundary intervals for pairs of distributions. This procedure is explained in detail by King in his discussion of the implementation of the interval preference measurement technique.

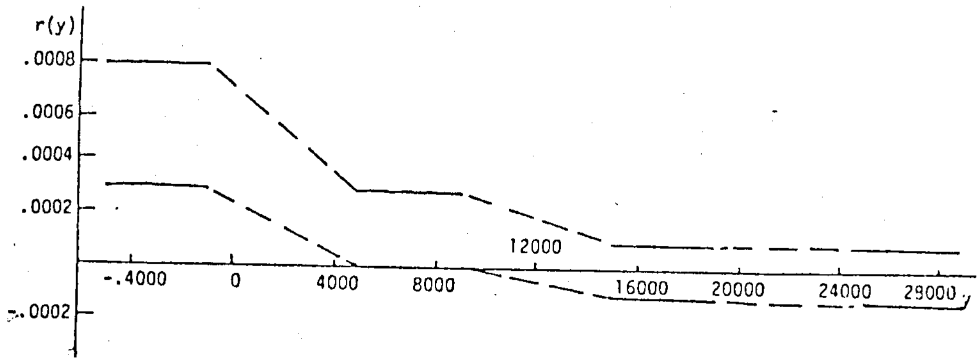

Interval Measurement for Decision Maker A

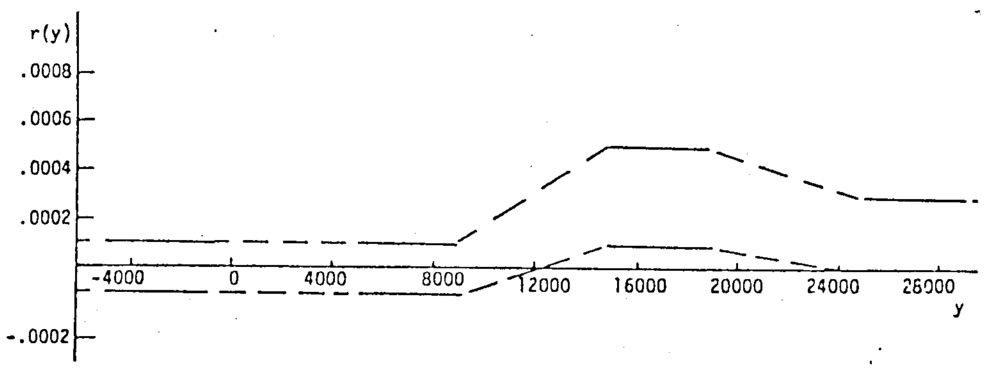

Interval Measurement for Decision Maker B

Figure 2. Interval Preference Measurements for Two Decision Makers

utility functions, on the other hand, their form is often severely limited by the functional form used to estimate the utility function (Lin and Chang). Finally, it should be noted that the interval approach to the measurement of preferences also avoids another common problem encountered in the estimation of single valued utility functions. Because all questions posed require a choice between two uncertain prospects, biases due to preference for or aversion to gambling per se (Officer and Halter) are avoided.

The greatest strength of this procedure, however, is its flexibility. Results of an experiment designed to test the predictive power of interval preference measurements and stochastic dominance with respect to a function against that of empirically estimated single valued utility functions and first and second degree stochastic dominance, which are summarized in Table 1, clearly show the problems associated with the more commonly used evaluative criteria. Single valued utility functions provide a complete ordering of choices, but they often exclude decision makers' preferred choices from the efficient set. Such errors can be likened to Type I errors in a statistical test. First and second degree stochastic dominance, on the other hand, rarely exclude a preferred choice from the efficient set. Often, however, they also fail to reduce the size of the efficient set. Such errors can be likened to Type II errors in a statistical test. Clearly, then, there are trade-offs between the accuracy and the discriminatory power of a preference measurement. The experimental results demonstrate that, unlike other measurement techniques and evaluative criteria, the combined use of stochastic dominance with respect to a function and interval preference measurements allows explicit consideration of these trade-offs. By altering the precision of interval preference measurements, the percent of choices predicted correctly and the percent of choices ordered were easily manipulated. As the precision of the interval measurements was increased, they became the basis for more discriminating preference orderings. The probability of excluding preferred choices from the efficient set also increased, however.

In addition to this experimental test, the interval approach to preference measurement has also been used in a more applied setting. It was used in a series of extension workshops as a tool for helping farmers think about their own risk attitudes.[1] Farmers found the choices to be interesting and had little difficulty in completing the questionnaire. The range of responses was quite broad, with individuals ranging from extremely risk averse to extremely risk loving. Several discernable patterns did emerge, however. Most decision makers exhibited increasing absolute risk aversion over lower income levels and decreasing absolute risk aversion at higher income levels. For most, the interval measurement of absolute risk aversion included negative as well as positive values at some level of income. In fact, only four of the seventeen decision makers for whom the questionnaires were analyzed had lower-level absolute risk aversion functions which were everywhere nonnegative. This casts serious doubt on the applicability of a criterion such as second degree stochastic dominance, which is valid only for decision makers who are risk averse at all income levels.

---

[1]The interval measurements shown in Figure 2 were made at one of these workshops.

Table 1. Performance Indicators for Alternative Preference Measures[a]

Performance Indicator	Interval Measurement: Number of Questions				Single Valued Utility Function	First Degree Stochastic Dominance	Second Degree Stochastic Dominance
	1	2	3	4			
1. Percent of choices predicted correctly[b]	98	88	78	72	65	100	98
2. Percent of choices ordered	9	50	83	91	100	0	7

[a]We thank Garth Carmen, who helped to conduct this experiment.

[b]A choice was said to be predicted correctly if the preferred distribution was not excluded from the efficient set.

A Generalized Computational Procedure for the Identification of Preferred Choices Under Uncertainty

Interval measurements of decision maker preferences provide information required to order any two specified choices. In most decision situations, however, a large if not infinite range of choices is open to the decision maker. As a result, some systematic technique for the identification and evaluation of a large number of possible strategies is also required for the implementation of stochastic dominance with respect to a function. Such a technique should be flexible enough to be applicable in a wide range of decision situations without requiring that important simplifying assumptions be made concerning decision maker preferences, the form of outcome distributions associated with the strategies considered, or the nature of the problem itself. Such a technique is introduced in this section.

Some of the shortcomings of mathematical programming models commonly used for the analysis of decisions made under uncertainty were identified in the introduction to this paper. They include unrealistic restrictions on decision maker preferences and probability distributions and limitations on the degree to which complex stochastic processes and flexible management strategies can be adequately considered. The risk efficient Monte Carlo programming (REMP) model developed by Anderson is, in many respects, an attractive alternative to other mathematical programming models. The REMP model employs Monte Carlo programming techniques (Donaldson and Webster) to construct a large number of feasible management strategies in a random fashion. The distribution of total net returns associated with each strategy under consideration is determined analytically under the assumption that distributions of net returns for each activity and distributions of total total net returns for each strategy are members of the beta family. The evaluative criterion of second degree stochastic dominance is used to evaluate strategies sequentially as they are generated. The REMP model allows for considerable flexibility in the representation of probability distributions, since the beta distribution can assume a variety of forms. It also places few restrictions on decision maker preferences, since second degree stochastic dominance requires only that decision makers be risk averse. As has already been noted, however, second degree stochastic dominance is not a very discriminating evaluative criterion, and efficient sets identified by the REMP model can be prohibitively large.

The generalized procedure for the identification of preferred choices described here is in many respects an extension of the REMP model. Feasible strategies are generated using a modified form of the Monte Carlo programming model that is the basis for Anderson's model. Under this more general procedure, however, a strategy can be defined by specific levels for choice variables, by a set of adaptive decision rules which use information from the environment to determine actions through time, or by some combination of the two. Probability distributions of outcome levels associated with the strategies considered are not determined analytically as in the REMP model. Rather, they are determined by simulating performance under each strategy for a large number of sample states of nature. The resulting sets of outcomes are used to define cumulative distribution functions for the outcome distributions associated with each strategy. This procedure facilitates the consideration of the impact of a wide range of random factors and allows for much greater flexibility in the representation of complex stochastic processes. Finally, strategies are evaluated using interval measurements of decision maker preferences in combination with the evaluative criterion of stochastic dominance with respect to a function.

Like the REMP model, this procedure is an iterative one. Strategies are generated and evaluated sequentially. The determination of a truly optimal choice is not ensured. If a sufficiently large number of plans is examined, however, it is reasonable to conclude that the efficient set will contain a nearly optimal choice. In applications to date, from 250 to 2,000 strategies have been evaluated. In practice the number considered depends on the complexity of the problem being analyzed, on the form of the management strategies, and on the perceived value of identifying strategies which are very nearly optimal.

Because of its similarity to the REMP model, this procedure for the identification of preferred choices under uncertainty can be called the generalized risk efficient Monte Carlo programming model (GREMP). Interrelationships among the three major processes within the model--strategy generation, outcome distribution determination, and evaluation--are illustrated in the flow chart in Figure 3. A more complete description of the model and a listing of the computer program which implements it are given in King.

Finally, several additional comments should be made about the implementation of the GREMP model. Computationally, it is relatively efficient. In the analysis of one test problem with 35 choice variables, 12 linear constraints, and a relatively simple simulation component, for example, 1,000 alternative strategies were generated and evaluated using less than 70 seconds of CPU time on a CDC 6500. Furthermore, the core size of the computer program which implements the GREMP model is relatively small, and the degree of computational accuracy required for internal calculations is not particularly great. This suggests that it may be possible to design software which will permit the use of the GREMP model on a moderately sized personal computer. With regard to ease of implementation, the computer program for the GREMP model can be easily adapted for use in the analysis of a wide range of problems. Several problem-specific subroutines must be supplied by the user, though, so some knowledge of computer programming is required.

An Example

The strengths of the GREMP model and of the interval preference measurements which are the primary source of normative information for its operation can, perhaps, be best illustrated with an example. The problem considered here involves the identification of a land rental and cropping strategy for a cash grain farmer who owns 240 tillable acres. Up to 320 additional acres can be rented, but land can be rented only in blocks of 80 acres. Corn and soybeans are the only crops grown. Yields for both crops are affected by growing conditions and by timeliness of planting and harvest, all of which are considered to be random factors. Product prices represent still another source of uncertainty in the problem.

A management strategy in this simple example is defined by four choice variables. Acres rented, v_1, must be an integer multiple of 80 and can range from 0 to 320. Acres of corn and soybeans, v_2 and v_3 respectively, are considered to be integer multiples of 10. Levels of these variables must conform to the following constraint:

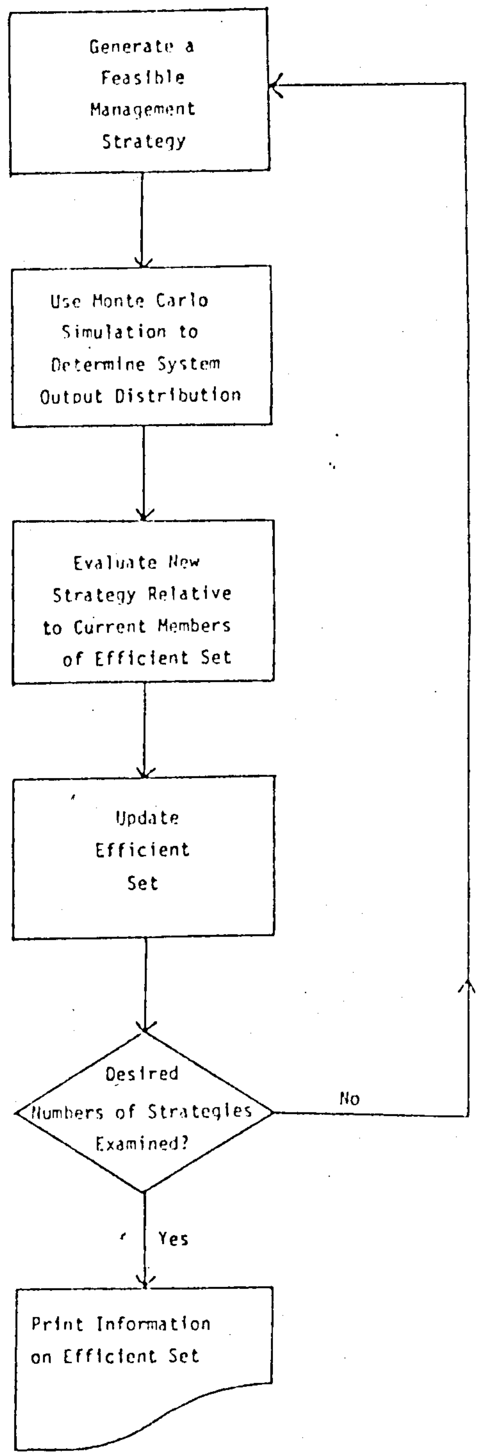

Figure 3. A Flow Chart of the GREMP Model

$$v_2 + v_3 = 240 + v_1,$$

which implies that the combined acreages of the two crops must equal total available acreage. Finally, because time available for field work is considered to be a random factor, it may be desireable in some states of nature to stop planting corn before the acreage level specified by v_2 is reached. This would be the case, for example, if poor weather in April and early May delayed planting to such an extent that expected corn yields would be unacceptably low. The fourth choice variable, v_4, then, is a date after which all unplanted acreage will be planted in soybeans regardless of the values of v_2 and v_3. Possible values for v_4 are: May 18, May 26, and June 3.

Subjective probability distributions for the three sets of random factors in this problem--crop prices, crop yields given specific planting and harvest dates, and time available for fieldwork during specific planting and harvest periods--were specified using techniques described in King, and Monte Carlo methods were used to construct 20 sample states of the environment.[1] A simple simulation model was specified to determine the level of net income realized under any particular management strategy in any given state of the environment. The simulation begins with the computation of charges for land rental, if any. Subject to time available for fieldwork, the model then simulates the planting of corn until the specified corn acreage is attained or until the date after which all remaining acreage is to be planted in soybeans. Planting of soybeans then proceeds until all acreage is planted or until June 19, the final day of the last planting period. There is no assurance that all available acres will be planted in a particular state of nature; this depends on levels of time available for fieldwork. Costs for seed, fertilizer, herbicides, and fuel are incurred only for acreage actually planted.

Harvesting is simulated in a similar manner. Subject to time available, soybeans are harvested as quickly as possible, with acreage planted first being harvested first. This continues until all planted soybean acreage is harvested or until the date is reached after which all unharvested acreage is judged to be a total loss. The harvest of corn then proceeds in a similar manner. Again, there is no assurance that all acres planted will be harvested. All harvested acreage is classified according to crop, planting period, and harvest period so that total production for each crop can be determined. Drying and hauling costs are assessed for each for each bushel harvested. Finally, receipts from crop sales are determined by multiplying the number of bushels of each crop harvested by the relevant crop price, and net income is computed by subtracting fixed and variable costs from this figure.

The GREMP model was used to identify an efficient set of choices for each of the two decision makers whose preference measurements are given in Figure 2. In each case 500 strategies are generated and evaluated. The efficient set of choices for decision maker A is comprised of the eight

---

[1] Crop price and days available for fieldwork distributions were considered to be multivariate beta, while crop yield distributions were considered to be multivariate normal. A procedure developed by King for the generation of variates from non-normal multivariate distributions was used to construct the sample environmental states.

strategies defined in Table 2. Levels of land rented range from 0 to 160, with four of the eight strategies calling for land rental levels of 80 acres. Soybeans are the predominant crop in each strategy, reflecting the fact that cost-price relationships favor soybeans in this example. Because of the low corn acreage levels, the switching rule is of little importance in these strategies. Mean income levels range from slightly less than $3,000 to slightly above $10,000. Minimum income levels vary little from one strategy to another, but maximum income levels are significantly affected by land rental values.

The efficient set of choices for decision maker B is comprised of the nine strategies defined in Table 3. In this case land rental levels tend to be higher than those called for in the strategies included in the efficient set of decision maker A. With the higher total acreage levels, the mix between corn and soybeans becomes more even, but most of the available acreage is still planted to soybeans in each strategy. Mean income levels tend to be higher in this set of strategies, but the dispersion of possible income levels is also greater. Given the differences in the preference measurements for the two decision makers, the dissimilarity between the two efficient sets is understandable. The interval measurement of absolute risk aversion for decision maker A indicates a high level of absolute risk aversion over negative income levels, which implies that he has a strong aversion to losses. Decision maker B, on the other hand, has much lower levels of absolute risk aversion at low income levels, and his efficient set contains strategies which provide opportunities for the realization of relatively high income levels but which can also result in substantial losses.

Several general comments can be made about these results. First, they provide clear evidence of the discriminatory power of interval preference measurements and stochastic dominance with respect to a function. Both efficient sets contain less than 2 percent of the total number of strategies examined. Second, these results demonstrate that preferences have an important impact on the choices made by individuals, which implies that explicit consideration should be given to them in an applied decision analysis. Finally, the results help to illustrate the power of the GREMP model. Though the problem considered here is a relatively simple one, it would be extremely difficult to solve, as formulated, using more conventional mathematical programming techniques due to the nature of the management strategy, which has integer choice variables and is flexible enough to allow for changes in crop mix in response to environmental conditions, and due to the complexity of the stochastic process that is modeled.

Implications for Future Research

The two related analytical techniques introduced in this paper were designed to facilitate the application of decision theory based on the expected utility hypothesis in a practical context. They help to resolve some of the problems that have limited the usefulness of this valuable theoretical tool in the solution of actual decision problems in which uncertainty is an important factor. Both techniques have been used in recently completed or ongoing research projects. The interval approach to preference measurement is currently being used in a study designed to test for correlations between decision maker attributes and their attitudes toward risk taking (Carman). This study will provide valuable additional experience with

Table 2. Efficient Strategies for Decision Maker A

Efficient Strategy	Control Variable Levels				Properties of Net Cash Income Distribution			
	Acres Rented	Acres Corn	Acres Soybeans	Switching Date	Mean	Standard Deviation	Minimum Value	Maximum Value
1	0	0	240	June 3	3816	8357	-11875	20368
2	160	120	280	May 26	10152	12517	-12468	30877
3	80	60	260	May 18	7138	10328	-12102	25972
4	160	110	290	May 26	9936	12526	-12605	31170
5	80	50	270	May 18	7168	10437	-12220	26419
6	80	80	240	June 3	6994	10142	-11865	24305
7	0	50	190	May 26	2949	7691	-11482	15963
8	80	70	250	June 3	7079	10239	-11983	25165

Table 3. Efficient Strategies for Decision Maker B

Efficient Strategy	Control Variable Levels				Properties of Net Cash Income Distribution			
	Acres Rented	Acres Corn	Acres Soybeans	Switching Date	Mean	Standard Deviation	Minimum Value	Maximum Value
1	160	160	240	May 26	9840	12199	-12987	27685
2	240	200	280	May 26	10808	15481	-17272	33952
3	160	120	280	May 26	10152	12517	-12168	30877
4	160	140	260	May 18	10167	12466	-12231	30331
5	240	190	290	May 26	10798	15483	-17272	33952
6	160	130	270	May 18	10193	12472	-12350	30366
7	160	140	260	June 3	10088	12395	-12231	29610
8	160	130	270	May 26	10167	12458	-12350	30366
9	160	150	250	June 3	9949	12367	-12604	28764

this procedure. The GREMP model was used in an extension of the example presented in the preceeding section to evaluate combined production and marketing strategies for an agricultural firm (King). Of particular interest in this application is the incorporation of adaptive forward contracting decision rules into a larger marketing strategy. These flexible strategies, the parameters of which are choice variables within the model, make forward contracting decision over a nine-month period dependent upon price expectations, price movements, and the degree to which the crop production plan has been successfully implemented. The GREMP model is also currently being used in a study concerned with investment and disinvestment decisions by electric utilities. Two possible future applications are in the areas of integrated pest management and on-farm water management. In the pest management study, the GREMP model will be used to identify flexible weed control strategies for two cropping systems in Colorado and to determine the value of weed infestation predictions based on weed seed counts made prior to planting. In the irrigation study, it will be used to evaluate alternative cropping and irrigation strategies for Egyptian farmers. In both of these studies it is hoped that interval preference measurements can be made for representative samples of decision makers.

REFERENCES

Anderson, J.R. "Programming for Efficient Planning Against Non-Normal Risk." Australian Journal of Agricultural Economics, 19(1975): 94-107.

Arrow, K. J. Essays in the Theory of Risk-Bearing. Chicago: Markham Publishing Company, 1971.

Boussard, J. M. and M. Petit. "Representation of Farmers' Behavior Under Uncertainty with a Focus-Loss Constraint." Journal of Farm Economics, 49(1967): 869-80.

Carman, G. "The Relationship of Producer Attributes to Risk Preference and Management." Unpublished Type III seminar paper, Department of Agricultural Economics, Michigan State University, 1979.

Cocks, K. D. "Discrete Stochastic Programming." Management Science, 15(1968): 72-79.

Donaldson, G. F. and J. P. G. Webster. "An Operating Procedure for Simulation Farm Planning--Monte Carlo Method." Department of Agricultural Economics, Wye College, University of London, 1968.

Hazell, P. B. R. "A Linear Alternative to Quadratic and Semivariance Programming for Farm Planning Under Uncertainty." American Journal of Agricultural Economics, 53(1971): 53-62.

King, R. P. "Operational Techniques for Applied Decision Analysis Under Uncertainty." Unpublished Ph.D. dissertation, Department of Agricultural Economics, Michigan State University, 1979.

Lin, W. W. and H. S. Chang. "Specification of Bernoullian Utility Function in Decision Analysis." Agricultural Economics Research, 30(1978): 30-36.

Massé, P. Optimal Investment Decisions. Englewood Cliffs, New Jersey: Prentice-Hall, Inc., 1962.

McInerney, J. P. "Linear Programming and Game Theory Models." Journal of Agricultural Economics, 20(1969): 269-78.

Meyer, J. "Choice Among Distributions." Journal of Economic Theory, 14(1977a): 326-36.

_______. "Second Degree Stochastic Dominance with Respect to a Function." International Economic Review, 18(1977b): 477-87.

Officer, R. R. and A. M. Halter. "Utility Analysis in a Practical Setting." American Journal of Agricultural Economics, 50(1968): 257-77.

Paris, Q. "Revenue and Cost Uncertainty, Generalized Mean-Variance, and Linear Complementarity Problem." American Journal of Agricultural Economics, 61(1979): 268-75.

Pratt, J. W. "Risk Aversion in the Small and in the Large." Econometrics, 32(1964): 122-36.

Rae, A. M. "Stochastic Programming, Utility, and Sequential Decision Problems in Farm Management." American Journal of Agricultural Economics, 53(1971): 448-60.

Robison, L. J. and R. P. King. "Specification of Micro Risk Models for Farm Management and Policy Research." Agricultural Economics Report No. 349, Department of Agricultural Economics, Michigan State University, 1978.