AEWP 1986-20

Risk Modeling via Direct Utility Maximization Using Numerical Quadrature

Michael S. Kaylen, Paul Preckel, and Edna Loehman*

*Michael Kaylen is Assistant Professor, Department of Agricultural Economics, University of Missouri-Columbia; Paul Preckel is Assistant Professor and Edna Loehman is Associate Professor, Department of Agricultural Economics, Purdue University.

Introduction

Uncertainty of future events has led to the development of a number of models to analyze choice in situations with risk. Since some information regarding the choices to be made is not known with certainty, the assumption is made that decision makers maximize expected utility. That is, an explicit subjective joint probability distribution is ascribed to the unknown variables (e.g,. the price and quantity of output for agricultural production). The expected value of utility serves as the measure by which the decision maker evaluates choices. When the probability distribution in question is continuous, the computation of expected utility amounts to integrating a function. Frequently, these integrals have no closed form solution. Because of a lack of computational procedures, choices have been ranked using approximations of expected utility [e.g., expected value - variance (E-V) models (Tobin, Markowitz)] or stochastic dominance analysis (Cochran, et al., Meyer).

The purpose of this article is to discuss a new approach to risk modeling. As with the direct expected utility maximizing nonlinear program (DEMP) approach recently proposed by Lambert and McCarl, the central idea is to use mathematical programming to directly solve the expected utility maximizing problem. The main difference is that the new approach explicitly recognizes that the objective function for this problem typically involves an integral. Whereas Lambert and McCarl obviate the necessity of evaluating the integral by advocating the use of an empirical density function, the technique proposed in this article uses quadrature methods from the numerical analysis literature to compute the integral. (Note that the DEMP approach may also be viewed as a quadrature technique, as will be discussed below.) For this

reason, we refer to our approach as the direct expected utility maximizing program using quadrature (DEMPQ).

This article procedes as follows. In the first section, numerical quadrature procedures are briefly discussed. The next section discusses the application of the quadrature routines to the utility maximization problem. The third section contains examples contrasting DEMP and DEMPQ. The final section consists of some concluding comments.

Numerical Quadrature Discussion

Consider the problem of minimizing the parametric integral:

$$\min_{T} F(T) = \int_{a_n}^{b_n} \int_{a_{n-1}}^{b_{n-1}} \cdots \int_{a_1}^{b_1} f(X;T)\, dX_1\, dX_2 \cdots dX_n \qquad (1)$$

where X is an n-dimensional vector of variables of integration, and T is a vector of choice variables. If the integral in (1) has no closed form solution, a quadrature routine must be used to perform numerical integration. That is, F(T) is approximated by I(T), a weighted sum of function values at selected points:

$$I(t) = \sum_{j=1}^{s} w_j\, f(X_j;T), \qquad (2)$$

where $\{w_j\}$ are the weights and $\{X_j\}$ are points in the integration region.

There are essentially two methods of choosing the X_j's to be used in (2). (See Haber for a survey of quadrature techniques). The first method is a Monte Carlo procedure, while the second is a systematic approach. The simplest version of the Monte Carlo approach is to generate a random sample from a population which is uniformally distributed over the integration region. Denoting the sampled points $\{X_j\}$, we have

$$I(T) = \sum_{j=1}^{s} R[f(X_j;T)/s],$$

where $R = \prod_{i=1}^{n} (b_i - a_i)$. The approximation to the integral is thus seen to be simply the volume of the integration region times the average value of the function evaluated at the sampled points.

For large s, the error in the corresponding integral estimate is distributed normally with zero mean and a standard deviation of $Rs^{-\frac{1}{2}}\ \sigma(f)$, where $\sigma(f)$ denotes the standard deviation of $f(X;T)$, treating X as being uniformally distributed on the integration region. While no deterministic error bound is available, the probability of an error outside a given range is seen to converge to zero at the rate $s^{-\frac{1}{2}}$.

Other Monte Carlo methods have been devised to try to improve upon the standard deviation of the error term. These involve more complex sampling procedures. For example, in "importance sampling", the idea is to construct a function g(X;T) such that

$$\int_{a_n}^{b_n} \int_{a_{n-1}}^{b_{n-1}} \cdots \int_{a_1}^{b_1} g(X;T)\ dX_1\ dX_2 \cdots dX_n = 1,$$

g ( ) is nonnegative, and g/f is approximately constant for all X. The X_j's in (2) are then determined by sampling from a population with a probability density function given by g. The integral estimate obtained using (2), with f replaced by f/g, then has an error with mean zero and a standard deviation of $Rs^{-\frac{1}{2}}\ \sigma(f/g)$. The decrease in the standard deviation of the error estimate is thus seen to depend on how closely g "mirrors" f.

Most systematic quadrature routines involve choosing the X_j's and w_j's in (2) in such a way as to ensure the integral approximation has some desirable property. A common approach is to ensure the integral approximation is exact for all polynomials of a specified degree.

A particularly simple quadrature method involves approximating the function by constants over intervals of the domain. The idea is to define a grid over the integration region by dividing the interval for each variable of integration into M equally-spaced parts. The formula is then given by:

$$I(T) = \sum_{j=1}^{M^n} R_M f(X_j)$$

where R_M is the volume of the rectangles defined by the grid $[M^{-n} \prod_{i=1}^{n} (b_i - a_i)]$ and the X_j's are the centers of the rectangles. This approximation is referred to as the M^n-point product repeated midpoint formula. It is illustrated in Figure 1 for n = 1 and M = 1, 2, 4.

The M^n-point product repeated midpoint formula may be performed using a sequence of values for M, say $M = M_1, M_2, \cdots, M_p$ $(M_1 < M_2 < \cdots < M_p)$. These approximations to the integral may then be combined by making use of an extrapolation technique due to Richardson and Gawnt. For every sequence, constants can be found such that the linear combination of the individual approximations, using these constants, will often be a much better approximation to the integral than any of the individual approximations. The constants are chosen to ensure this formula will be exact for functions which are polynomials of degree less than or equal to 2P-1 (Dahlquist and Björck).

Kaylen and Preckel have recently developed a FORTRAN subroutine which can be used to approximate the integral in (1) and its derivatives with respect to the parameters, T. The subroutine is called MINTDF for Multiple INTegral - Differentiable Function. It uses Richardon's extrapolation, and either of two sequences for M may be used:

$$M = M_0, 2M_0, \ldots, PM_0, \text{ or} \tag{3}$$

$$M = 2^0 M_0, 2^1 M_0, \ldots, 2^{p-1} M_0. \tag{4}$$

The first sequence was analyzed by Lyness and McHugh and is referred to as the

progressive procedure. An obvious advantage to it is that potentially far fewer function evaluations are required for a given P. However, the latter sequence has the advantage that it is guaranteed to converge for all Riemann - integrable functions. For this reason, all results reported in this paper use (4), with $M_0 = 1$.

To a large extent, the choice of quadrature technique depends upon the context within which it is to be used. In the case at hand, the solution of (1) will typically involve evaluating F( ) at many points. Even if F( ) is continuous in T, I( ) may very well have discontinuities unless it is evaluated using the same set of X's and the same extrapolations. These discontinuities can seriously hamper the ability of the optimization program to find an acceptable solution (see Gill, et al., Ch. 7). For this reason, the reader is cautioned against using unmodified library routines for numerical integration without a close examination of their workings. Using a Monte-Carlo type integration routine would entail saving the points at which (2) is evaluated, potentially causing very large memory requirements. On the other hand, the systematic approach can, at a very low computational cost, simply recalculate the set of points.

For many-dimensional problems, the convergence rate (in a probablistic sense) for the Monte Carlo procedure is faster than that for the deterministic approaches (on the order of $s^{-1/n}$, where s is the number of points at which the function is evaluated and n is the number of dimensions). This advantage of the Monte Carlo approach is somewhat counterbalanced by the fact that the systematic approaches give deterministic, as opposed to probablistic, error measures. It should be noted that the Monte Carlo procedure also involves the extra computational cost of sampling from a distribution.

Quadrature and the Expected Utility Problem

The form of the expected utility maximization problem discussed in this paper is given by

$$\underset{T}{\text{Max}}\ E[u(W)] \tag{5}$$

$$\text{subject to } g(T) \leq b,\quad T \geq 0.$$

T is an r - dimensional vector of choice variables, E is the expectation operator, u(W) is utility as a function of wealth, W = f(T;X), X is an n - dimensional vector of random variables, and $g(T) \leq b$ represents q arbitrary constraints on the choice set (which may involve integrals as in chance - constrained problems). Assuming X has a joint probability density function given by $p(X_1, \ldots, X_n)$, the objective function may be written as

$$\underset{T}{\text{Max}} \int_{a_n}^{b_n} \cdots \int_{a_1}^{b_1} u[f(T;X)]p(X_1,\ldots, X_n)dX_1 \ldots dX_n$$

where the domain of X_i is $[a_i, b_i]$, $i = 1, \ldots, n$.

The subroutine MINTDF requires the a_i's and b_i's to be finite numbers. This may present a problem for some probability distributions. One ad hoc solution is to use MINTDF to evaluate the integral of the probability density function for increasing ranges until the calculated area stabilizies at one. For example, in this paper a normal distribution is used. MINTDF was used to compute the area under a standard normal probability density function using $(a_1, b_1) = (-k, k)$ for various levels of k and an error bound of 10^{-14}. The calculated area was seen to stabilize at k = 8, suggesting a range of ± eight standard deviations about the mean might be acceptable for normally distributed variables.

One of the strengths of MINTDF is the ability to specify an error bound. The question is how best to use it. For complicated problems, the solution of

(5) may require very many function evaluations (trial values for T). If the sequence (4) is used, this implies the argument of the integral will be evaluated very many times. Consequently, we suggest an initial "rough" solution be obtained by specifying a fairly small value for P in (4). A "refined" solution may then be obtained using the "rough" solution as the initial solution and, rather than specifying P, bounding the error at some acceptable level. For example, for the problems in this paper, we specified P = 5 for the "rough" solution, then set an error bound equal to one percent of the "rough" integral estimate. Since the user is probably most interested in the solutions for the choice variables, it may be advisable to check if the "refined" solution differs dramatically from the "rough" solution. If it does, an even lower error bound may be advisable.

Model Experimentation

The above discussion has focused on two available techniques for implementing direct utility maximization problems. This section will address the differences between these methods by examining their performance on some simple examples. The models used for the experiments reported here are similar to those used to demonstrate the DEMP approach (Lambert and McCarl).

The model employed for this example may be described as a two asset portfolio model wherein the risk averse agent wishes to maximize the expected utility of wealth subject to a fixed endowment of capital. In mathematical terms the model is:

$$\max_{T} E\,[u(W)]$$

$$\text{subject to } W = X_1 T + X_2 (1-T),$$

$$0 \leq T \leq 1.$$

where X_1 and X_2 are random variables representing the yields associated with

the two alternative investments. Two utility functions, which represent those most commonly used for empirical work, will be considered -- namely:

the quadratic -- $u(W) = 100\,W - W^2$, and

the negative exponential -- $u(W) = -e^{-W}$.

The scaling of the random yields was chosen so as to scale wealth to ensure the risk aversion of the agent is significant. (Note that the selection of a risk aversion coefficient for these two utility functions is equivalent to rescaling wealth.)

Three types of distributional assumptions for the random variables X_1 and X_2 are made: normal, symmetric triangular, and uniform. Due to the lack of availability of a reliable multivariate random number generator for correlated random variables, X_1 and X_2 were treated as independent random variables. While this facilitates the comparison between the two alternative approaches to the integration problem, this assumption could be relaxed for either method without additional difficulties. (Note however, that if the sample of points for the region of integration must be generated, as in the DEMP or Mont-Carlo type approach, then significant care is needed in designing the sampling procedure.)

A description of the distributions used in this experiment is displayed in Table 1. For purposes of implementing the DEMP or Monte-Carlo type integration procedure, a sample of one thousand points was drawn from each of the distributions for each of the two types of utility functions. Descriptive statistics for each of the samples are also displayed in Table 1. It appears that the parameters of the distributions and the associated sample statistics are substantially in agreement, suggesting that reasonable "draws" from the distributions were obtained.

One advantage of the test problems created for this experiment is that they are sufficiently simple that the exact analytic solutions may be derived. These solutions are displayed in the first two lines of Table 2. (Note that the actual optimal values for T are the same for all of the instances where the quadratic utility function is used. This is consistent with the fact that when the utility function is quadratic, only the mean and variance are of importance to the optimization. It would be more efficient to solve these cases using quadratic programming.)

The next two lines in the table report the solutions computed when the Monte-Carlo type integration procedure was used (labeled DEMP). These solutions were based on the samples over the integration region that were reported in Table 1. Note that the percentage error in the calculation of the choice variable, T, is fairly large. It is greater than 18% on average, and in one case, it is greater than 40%. Likewise, the estimates of expected utility calculated at the true optimum value for T, are also rather inaccurate -- on average they are in error by 31%, and in one case the error is over 99%. This integration procedure seems to have the greatest difficulty with the negative exponential utility function.

For the DEMPQ procedure, two termination criteria were applied. The solutions labeled DEMPQ-1 were computed using 341 points (5 refinement levels), approximately one third the number of points used by DEMP. Even so, the average error in the calculation of T was less than 1.5%, and never greater than 4%. Similarly, the error in the estimates of expected utility were small -- less than 1.5% on average, and no greater than 6%. The solutions labeled DEMPQ-2 employ a different termination criteria for the integral. At each refinement level beyond the first, an estimate of the inaccuracy of the integral is available. A feature of the quadrature

procedure employed is the opportunity to specify the desired level of estimated accuracy. For this second set of DEMPQ solutions, the accuracy level was set to one percent of the level of expected utility computed for the DEMPQ-1 solutions. The percentage error in the calculation of T for these cases was less than 0.1% on average, and never greater than 0.5% The error in the computation of expected utility was also less than 0.1% on average, and never greater than 0.5%. The cost of this greater accuracy is a greater number of points being evaluated (in some cases as many as five times the number used by DEMP). In the case where the utility function was quadratic and the distribution function was uniform, the number of points evaluated was very small -- 21. That is consistent with the fact that only a limited number of refinements needs to be performed when the integrand is a polynomial. This is also the explanation for the small number of points that were needed for the case where the utility function was quadratic and the distribution function was triangular -- implying that the integrand was piecewise-cubic.

Concluding Remarks

Direct expected utility maxmization appears to be not only feasible but practical given the current state of the art for nonlinear programming software. This seems to be true even when the integrals involved in taking the expectation have no closed form solution. This fact opens new avenues for experimentation with nontraditional utility and distribution functions.

In this paper, two alternative methods for implementing such models have been compared. The primary issue seems to be what method is appropriate for performing the numerical integration. There are essentially two practical methods available: Monte-Carlo type (DEMP) and deterministic type (DEMPQ) methods. The Monte-Carlo type method has the advantage of being easy to

implement. It may be particularly preferred when the distribution is to be approximated by observed data directly (as opposed to generating a sample from an estimated distribution). This method has the disadvantage that a very large number of sample points are needed to accurately approximate the integral, and there is no apparent easy method for estimating the accuracy of the integral or the problem solution. The deterministic type method has the advantage that an estimate of the accuracy of the integral is readily available. This type of method is less easy to implement, and requires the estimation of the joint distribution function of the variables of integration.

Future work in this area will focus on the choice of appropriate distribution and utility functions, and the consequences of approximating the expected utility maximization problems with E-V models. Another area of interest is the use of chance constraints and safety-first type models. That is, it is not necessary to restrict the functions involving nonanalytic integrals to the objective function. It is quite feasible to incorporate these functions directly into the constraint set. Tools for numerical integration should eliminate the need for, and help evaluate the consequences of, making approximations to problems involving risk.

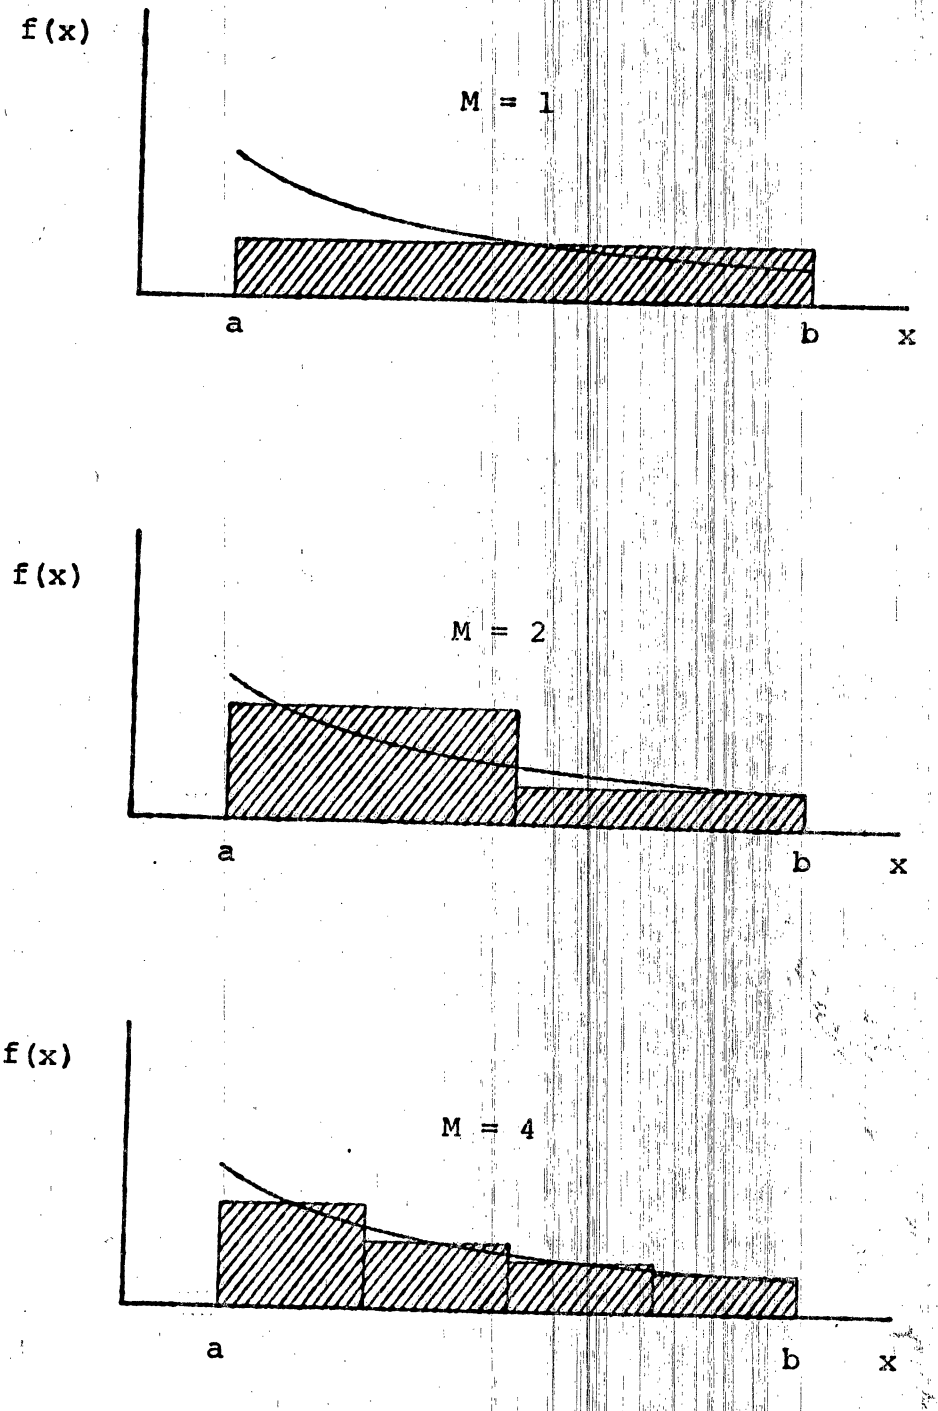

Note: Shaded areas represent estimates of the areas below f(x) between x=a and x=b.

Figure 1. Illustration of M^n - point product repeated midpoint formula.

Table 1. Description of Actual Distributions and Statistics for the samples Used by DEMP.

	Normal		Triangular		Uniform	
	X_1	X_2	X_1	X_2	X_1	X_2
Actual						
Min. Value	-	-	-24.64	-7.49	-14.49	-2.95
Max. Value	-	-	44.64	23.49	34.49	18.95
Mean	10	8	10	8	10	8
Variance	200	40	200	40	200	40
Correlation	0		0		0	
Generated for Quadratic Utility						
Min. Value	-37.78	-9.04	-22.68	-6.68	-14.37	-2.94
Max. Value	58.54	31.19	42.88	23.10	34.46	18.94
Mean	9.32	7.91	10.36	7.97	9.46	7.98
Variance	193.45	43.27	202.91	41.26	195.67	38.75
Correlation	0.05		-0.02		-0.02	
Generated for Neg. Exp. Utility						
Min. Value	-32.35	-10.29	-22.83	-6.13	-14.30	-2.95
Max. Value	47.60	27.58	43.77	22.71	34.49	18.93
Mean	9.72	7.91	9.93	8.01	9.94	8.27
Variance	209.78	41.73	206.35	38.70	192.10	39.00
Correlation	0.01		0.01		0.02	

Table 2. Analytic and Computed Solutions.

	Quadratic			Negative Exponential		
	Normal	Triangular	Uniform	Normal	Triangular	Uniform
Analytic Solutions						
T x 100	50.82	50.82	50.82	17.50	8.75	7.02
E(u)	759.02	759.02	759.02	-4125.74	-3.97	-0.60
DEMP Solutions [a]						
T x 100	42.76	56.71	42.89	17.83	12.34	8.89
E(u)	727.52	773.57	739.48	-5.69	-1.70	-0.47
DEMPQ-1 Solutions [a]						
T x 100	52.60	50.82	50.82	17.91	8.94	7.01
E(u)	753.47	759.02	759.02	-4176.93	-4.18	-0.59
E(u) Error Estimate [b]	235.91	0.93	0.00	1074.69	1.17	0.04
DEMPQ-2 Solutions [a]						
T x 100	50.68	50.82	50.82	17.48	8.75	7.02
E(u)	758.37	759.02	759.02	-4114.56	-3.97	-0.60
E(u) Error Estimate [b]	4.98	0.93	0.00	0.12	0.01	0.00
Number of Points	1365	341	21	5461	5461	1365

[a] DEMP and DEMPQ-1 used 1000 and 341 points, respectively, to evaluate the integral. The number of points for DEMPQ-2 was determined dynamically so that the error estimate would be no larger than one percent of the optimal expected utility computed by DEMPQ-1.

[b] The expected utility error estimate refers to the error computed by DEMPQ in the integral evaluation process.

REFERENCES

Cochran, Mark J., Lindon J. Robison, and Weldon Lodwick. "Improving the Efficiency of Stochastic Dominance Techniques Using Convex Set Stochastic Dominance." American Journal of Agricultural Economics 67(1985): 289-295.

Dahlquist, Germund and Åke Björck. Numerical Methods. Englewood Cliffs, New Jersey: Prentice Hall, 1974.

Gill, Philip E., Walter Murray, and Margaret H. Wright. Practical Optimization. New York: Academic Press, 1981.

Haber, Seymour. "Numerical Evaluation of Multiple Integrals." SIAM Review 12 (1970): 481-526.

Kaylen, Michael S. and Paul V. Preckel. "MINTDF User's Guide." Agricultural Economics Working Paper No. 1986-22. University of Missouri-Columbia.

Lambert, David K. and Bruce A. McCarl. "Risk Modeling Using Direct Solution of Nonlinear Approximations of the Utility Function." American Journal of Agricultural Economics 67(1985): 846-52.

Lyness, J.N. and B.J. McHugh. "Integration Over Multidimensional Hypercubes: I. A Progressive Procedure." Computer Journal 6 (1963): 264-70.

Markowitz, H. Portfolio Selection: Efficient Diversification of Investments. Cowles Foundation Monograph No. 16. New York: John Wiley & Sons, 1959.

Meyer, Jack. "Choice among Distributions." Journal of Economic Theory 14 (1977): 326-36.

Richardson, Lewis F. and J. Arthur Gawnt. "The Deferred Approach to the Limit." Philosophical Transactions of the Royal Society (A) 226 (1926): 299-361.

Tobin, J. "Liquidity Preference as Behavior towards Risk." Review of Economic Studies 25 (1958): 65-86.