The University of Western Australia

AGRICULTURAL ECONOMICS

SCHOOL OF AGRICULTURE

MODELLING RISK IN LINEAR PROGRAMMING USING DIRECT SOLUTION OF LINEARLY SEGMENTED APPROXIMATIONS OF THE UTILITY FUNCTION

David J PANNELL*

Agricultural Economics
Discussion Paper 2/88

Nedlands, Western Australia 6009

MODELLING RISK IN LINEAR PROGRAMMING USING DIRECT SOLUTION OF LINEARLY SEGMENTED APPROXIMATIONS OF THE UTILITY FUNCTION

David J PANNELL*

Agricultural Economics
Discussion Paper 2/88

The University of Western Australia
and
WA Department of Agriculture, South Perth, 6151 Western Australia

* The author wishes to thank Ross Kingwell, King-Min Wang, Bob Lindner, Rob Fraser, Joyotee Smith, Juliann Lloyd-Smith, Dave Falconer, Bruce Mattinson and Chris Arnold for helpful comments and fruitful discussion.

MODELLING RISK IN LINEAR PROGRAMMING USING DIRECT SOLUTION OF LINEARLY SEGMENTED APPROXIMATIONS OF THE UTILITY FUNCTION.

There are a number of techniques for representing attitudes to risk in mathematical programming models. Most are either inconsistent with expected utility maximisation (the predominant economic paradigm for risk analysis; see Hey 1979, 1981) or only maximise expected utility given a set of strong, restrictive assumptions. For example, Anderson et al. (1977) criticise the game theory (McInerney 1967) and maximum admissible loss (Boussard 1971) approaches for their poor basis in traditional risk theory, preferring an expected value - variance (E-V) approach using quadratic programming or its linear programming (LP) approximation, MOTAD (Hazell 1971). However, Lambert and McCarl (1985) summarise a number of criticisms of the assumptions underlying E-V analysis.

These criticisms are that E-V analysis assumes at least one of the following:

(a) a quadratic utility function, which implies increasing risk aversion with increasing wealth (Arrow 1971);

(b) an underlying normal distribution of wealth, whereas the real world may be characterised by asymmetric distributions (Hanoch and Levy 1969);

(c) small risk relative to wealth; and

(d) that E-V solutions are reasonable approximations of expected utility maximising solutions (Hanoch and Levy 1969).

Lambert and McCarl (1985) describe a method of risk programming using a direct expected utility maximising nonlinear program (DEMP). Their formulation is completely consistent with traditional risk theory and overcomes the main theoretical criticisms of E-V analysis.

DEMP has the desirable characteristics that it is free of restrictions on the form of the utility function (including whether it is concave or convex - it need only be quasi-concave) and that it is free of assumptions regarding the distribution of uncertain parameters. They note that "if the utility function were concave, one could . . . solve the problem as a linear program". This paper pursues that suggestion and presents a linear programming formulation (DELP for direct expected utility maximising linear program) with all the theoretical attractions of Lambert and McCarl's model except that it requires a concave utility function. Since this implies risk aversion, it is a tolerable restriction.

In the next section, the model is presented formally. Following this a simple numeric example is given and used to illustrate output interpretation. The accuracy of the DELP approximation relative to DEMP is tested. The paper concludes with some comments on the strengths and limitations of DELP.

Problem Formulation

The LP problem is an approximation of Formulation B in Lambert and McCarl (1985, p.48) which is:

$$\underset{\underline{x}}{\text{Max}} \sum_{k=1}^{n} p_k \, U(w_k) \text{ Subject to } \underline{A}\ \underline{x} \le \underline{b}$$

$$- \underline{c}_k' \, \underline{x} + w_k = w_o \ , \ k = 1 \ . \ . \ . \ n$$

$$\underline{x} \ge \underline{0} \ , \ \underline{w}_k \ge \underline{0} \ .$$

where there are n possible states of nature, p_k is the probability of the k^{th} state of nature, $\underline{x}$ is a vector of decision variables, w_k is the incremental wealth in state k given a particular $\underline{x}$, $U(w_k)$ is the utility obtained in state k given $\underline{x}$, $\underline{A}$ is the matrix of technical coefficients, $\underline{b}$ is the vector of constraint limits, $\underline{c}_k$ is the vector of net wealth contributions per unit of $\underline{x}$ in the k^{th} state of nature, and w_o is initial wealth.

The approximation works as follows. The utility function $U(w_k)$ is approximated by linear segments as shown in Figure 1 and represented in LP using the method formulated by Duloy and Norton (1975). The LP submatrix for the function in Figure 1 is shown in Table 1. Each linear segment requires an activity, U_i, and the function as a whole requires one extra constraint, I. The coefficient of each U_i in constraint C is the wealth value at one of the corner points in

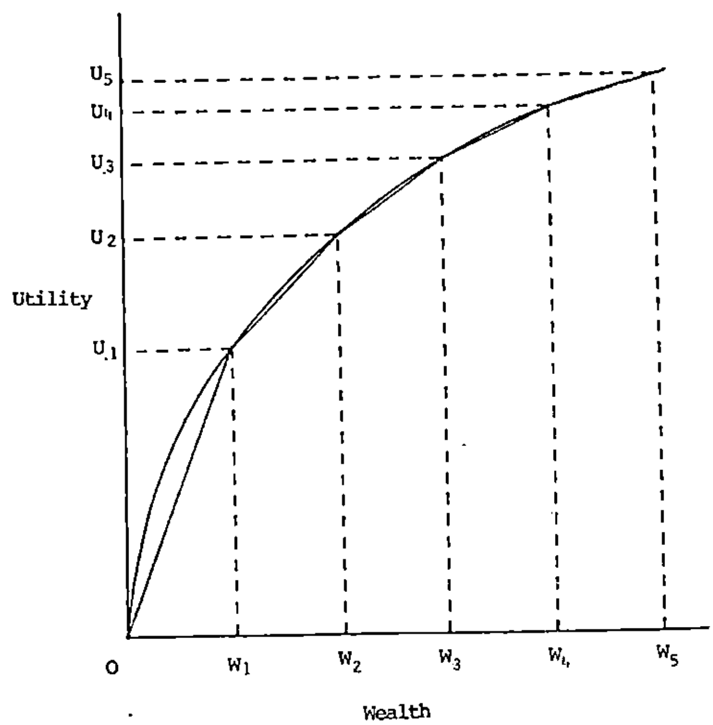

Figure 1: Linear Segmentation to Approximate Utility Function

Table 1: Approximation of Concave Utility Function as LP Objective Function

Name	x	U_1	U_2	U_3	U_4	U_5	B
C	-c	w_1	w_2	w_3	w_4	w_5	< = 0
I		1	1	1	1	1	< = 1
Obj (max)		u_1	u_2	u_3	u_4	u_5	

Figure 1, and the objective function value is the corresponding utility value. The I constraint limits selection of the U_i activities such that utility lies on or below the linearly segmented utility function in Figure 1. Efficiency requires that the solution lie on and not below the (approximated) utility function. Thus the solution will contain either one or two adjacent U_i activities. If the wealth level is between, say, w_2 and w_3, the objective function would be maximised by a combination of U_2 and U_2. Because the utility function is (assumed to be) concave, a combination of any other activities such as U_1 and U_3 would contribute less to the objective.

Whereas the non-linear programming formulation includes one activity for each state of nature, the LP formulation has one approximated utility function per state.

Thus the formulation of DELP is:

$$\max_{\underline{x}} \sum_{k=1}^{n} p_k \, \underline{u}_k' \, \underline{U}_k$$

Subject to $\underline{A}\ \underline{x} \leq \underline{b}$

$$-\ \underline{c}_k'\ \underline{x} + \underline{w}_k' \underline{U}_k \leq w_o \ , \ k = 1 \ . \ . \ . \ n$$

$$\underline{i}' \underline{U}_k \leq 1 \ , \ k = 1 \ . \ . \ . \ n$$

$$\underline{x} \geq \underline{0} \ , \ \underline{U}_k \geq \underline{0} \ .$$

where $\underline{U}_k$ is a vector of variables (one for each of m segments of the utility function) for the k^{th} state, $\underline{u}_k$ and $\underline{w}_k$ are m × 1 vectors of utility and wealth coefficients respectively (as per Table 1) for the k^{th} state, and $\underline{i}$ is an m × 1 vector of ones.

The following points should be noted about this formulation. As with MOTAD (Hazell 1971), a set of return vectors (one for each state of nature) is used to represent the variability in the system. The returns may be actual observations from recent years, or they may be subjectively determined. The probability of each state of nature may be 1/n if actual observations are used, or it may be subjective. The vector of probabilities can be scaled by any factor without affecting the optimal solution, since we are using ordinal utility functions which are unique only up to a linear transformation (Hey 1979, p.36). The returns and the probabilities together completely specify the variances and covariances and any other moments in the underlying multivariant distribution. Since data are entered directly, there is no need to calculate any moments.

As has already been noted, the utility function must be concave to avoid the problem described on pages 108-09 of Beneke and Winterboer (1973).

The segments used to approximate utility need not be the same for each state of nature. This allows the approximation to be made arbitrarily accurate by focusing the linear segments into narrower ranges in response to the initial solution.

An Example

I now present a simple example DELP model to illustrate the technique and to aid in discussing output interpretation. The example is taken from Anderson et al. (1977, ch. 7) where it is used to illustrate quadratic programming and MOTAD.

The problem is to choose between three crops (Standard wheat, Oats and New wheat) given four constraints (Total land, Wheat land, Capital and Labour). Resource requirements for each activity and resource availabilities are shown in Table 2. Activity net revenues for five sample years (states of nature) are to be used as the basis for planning next year's crops (Table 3).

Table 2: Resource Requirements and Availabilities for Example

	Standard wheat (ha)	Oats (ha)	New wheat (ha)	Resource availability
Total land (ha)	1	1	1	12
Wheat land (ha)	1	0	1	8
Capital ($)	30	20	40	400
Labour (man days)	5	5	8	80

Table 3: Sample of Activity Net Revenues ($/ha)

Prior year	Standard wheat	Oats	New wheat
1	99.8	68.3	112.7
2	133.3	130.4	238.4
3	142.7	33.3	93.9
4	154.3	74.4	83.2
5	11.4	25.4	109.7

The objective is to select the combination of crops which will maximise the farmer's expected utility given 'that in this case the farmer's preferences for profit over the relevant range can be adequately represented by the quadratic utility function $U = z - 0.0005z^2$' (Anderson et al. 1977, p.202).

For the purposes of this exercise, assume that initial wealth is zero. A preliminary step in constructing the DELP matrix is to select, for each state of nature, a number of levels of total net revenues to be used in approximating the utility function. The selection of points is arbitrary but it is important not to bound utility by selecting too narrow a range of revenues, so in the first instance they should encompass the largest and smallest possible net revenue totals for each state of nature. By inspection, it appears that total net revenue in any year will not exceed $1250 except in year 2, when it might be as high as $2500. Let us approximate utility by five linear segments; the 'corner' values are $250, $500, $750, $1000 and $1250 for most years, but $1500, $1750, $2000, $2250, $2500 for year 2. In both cases the point $0 is also the end point of a segment, but because its utility value is zero, a separate activity for $0 is not needed. The utility values for each of these points is calculated using the given utility function $U = z-0.0005z^2$.

For example, the first set of revenues gives utilities shown in Table 4. After an initial solution is found, smaller increments about the actually selected revenue levels can be used to improve accuracy. We will conduct a repeat of the analysis with increments of $50.

Table 4: Utility Levels for Various Wealth Levels

Wealth ($)	250	500	750	1000	1250
Utility	218.75	375.00	468.75	500	468.75

Part of the initial matrix for this problem is shown in Table 5. The first four constraints limit resource usage. Each of the next five, C1 to C5, transfers incremental wealth to the utility function for each state. Only the utility function for the first state of nature is shown in Table 5. The matrix would include four more similar sets of activities. The C constraints in this model are specified as equalities because the utility function used has negative marginal utility in parts of the relevant range. Realistic functions would be increasing everywhere and could be represented with 'less than or equal to' constraints. The constraints I1 to I5 restrain selection of the U_{km} activities such that any level of incremental wealth yields utility on or below the utility function.

Output Interpretation

Results are presented for solution of the model using the microcomputer package GULP (Table 6). Note that the objective function value, all shadow costs and shadow prices are expressed in utility terms and that the objective function is a linear transformation of the original utility function :

$$U = 5z-0.0025z^2.$$

Table 5: Part of LP matrix for DELP

Name	Standard-W	Oats (ha)	New-W	U11 (ha)	U12 (ha)	U13	U14	U15	B
Total land (ha)	1	1	1						<= 12
Wheat land (ha)	1		1						<= 8
Capital ($)	30	20	40						<=400
Labour (man days)	5	5	8						<= 80
C1 ($)	- 99.8	- 68.3	-112.7	250	500	750	1000	1250	= 0
C2 ($)	-133.3	-130.4	-238.4						= 0
C3 ($)	-142.7	- 33.3	- 93.9					...	= 0
C4 ($)	-154.3	- 74.4	- 83.2						= 0
C5 ($)	- 11.4	- 25.4	-109.7						= 0
I1				1	1	1	1	1	<= 1
I2									<= 1
I3								...	<= 1
I4									<= 1
I5									<= 1
Objective (utils)				218.75	375.00	468.75	500.00	468.75	

Table 6: Output for matrix in Table 4

Activity	Level	Shadow cost
Standard-W	3.81	0.00
Oats	0.00	18.15
New-W	4.16	0.00
U11	0.00	187.50
U12	0.00	62.50
U13	0.60	0.00
U14	0.40	0.00
U15	0.00	62.50
U21	1.00	0.00
U22	0.00	50.66
U23	0.00	163.83
U24	0.00	339.49
U25	0.00	577.65
U31	0.00	187.50
U32	0.00	62.50
U33	0.26	0.00
U34	0.74	0.00
U35	0.00	62.50
U41	0.00	187.50
U42	0.00	62.50
U43	0.26	0.00
U44	0.74	0.00
U45	0.00	62.50
U51	0.00	9.34
U52	1.00	0.00
U53	0.00	53.16
U54	0.00	168.81
U55	0.00	346.97

Constraint	Slack	Shadow price
Total land	4.03	0.00
Wheat land	0.03	0.00
Capital	119.24	0.00
Labour	27.66	0.00
C1	0.00	0.12
C2	0.00	-0.42
C3	0.00	0.12
C4	0.00	0.12
C5	0.00	0.58
I1	0.00	375.00
I2	0.00	1008.52
I3	0.00	375.00
I4	0.00	375.00
I5	0.00	81.19
Objective	2214.71	0.00

The optimal crop areas are 3.81 ha of standard wheat and 4.16 of new wheat. The shadow cost of oats is 18.15 utils. This value is difficult to relate to the original problem because it depends on the returns of three activities in five seasons and because it is an ordinal utility value. However, shadow costs are still useful for indicating the relative nearness of non-basic variables to entering the basis.

In interpreting levels of utility activities, it is helpful to recall the points selected as corners in the approximation of the utility function (Table 4). The solution includes 0.60 units of U13 and 0.40 of U14. Actual wealth for state 1 can be calculated as (0.60 × $750) + (0.40 × $1000) = $850. Similarly utility in state 1 is (0.60 × 468.75) + (0.40 × 500) = 480 (using the original utility function. This is also the contribution of state 1 to the expected utility total using the transformed utility function).

The shadow cost of utility activity, U_{kj}, indicates the vertical distance between the utility function of wealth W_j and the function's tangent at W*, the optimal wealth level in state k. For example, the shadow cost of U11 is the distance AB in Figure 2.

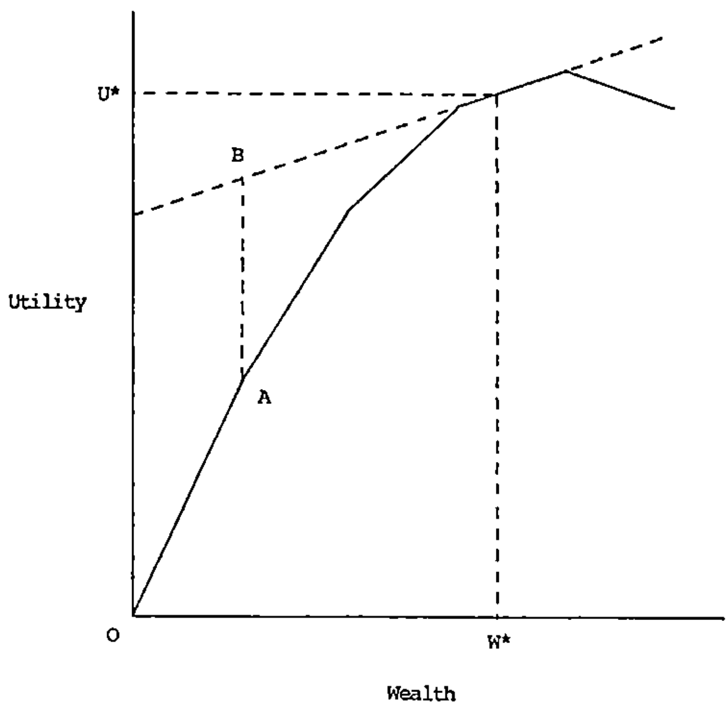

Figure 2: Utility Function for State 1

Output for resource constraints has its usual interpretation except that shadow prices are again in utility terms.

Slack for the constraints 1 to C5 must be zero as all wealth generated must be transferred to the utility function. That is why they are equality constraints in this problem; if C2 were a less than constraint, it would have positive slack rather than cause negative marginal utility. Clearly if a more realistic utility function were used, the wealth transfer constraints could be specified as 'less than or equal to' inequalities.

The equation for the tangent to the utility function in each state is given by the shadow prices of the C_k and I_k constraints which are, respectively, the tangent slope (marginal utility) and intercept for state k.

Slack for the I_k constraints will be zero unless actual wealth is between zero and w_{k1}, the first corner point.

Accuracy of the Approximation

In theory, DELP can be made an arbitrarily accurate approximation of DEMP by increasing the number of linear segments used to model utility or by focusing the linear segments into narrower ranges in response to an initial solution. In practice, however, it is important that acceptable accuracy be achieved without swamping the model with utility activities and without numerous iterations of focusing and re-solving. As a test of the practicality of obtaining accurate solutions from DELP, MINOS is used to solve the above example problem in a DEMP formulation. Results are compared with those from DELP using five linear segments and one round of focusing the segments (into wealth increments of $50). The original utility function implies extreme risk aversion, so the exercise is repeated with the squared term parameter reduced to -0.0004, -0.0003 and -0.0002.

Selected cropping activities and calculated total utilities are shown in Table 7.

Table 7: Accuracy of the DELP Approximation

Utility function	Item	DEMP	DELP	Percentage error
$u = z - .0005z^2$	Utility	2240.9	2240.5	0.0
	Standard wheat	3.79	3.82	+0.8
	Oats	0.00	0.00	0.0
	New wheat	3.81	3.74	-1.8
$u = z - .0004z^2$	Utility	2763.6	2762.0	-0.1
	Standard wheat	3.90	3.83	-1.8
	Oats	1.63	1.50	-8.0
	New wheat	4.10	4.17	+1.7
$u = z - .0003z^2$	Utility	3540.4	3540.0	0.0
	Standard wheat	3.33	3.35	+0.6
	Oats	4.00	4.00	0.0
	New wheat	4.67	4.65	-0.4
$u = z - .0002z^2$	Utility	4438.2	4426.5	-0.3
	Standard wheat	1.33	1.33	0.0
	Oats	4.00	4.00	0.0
	New wheat	6.67	6.67	0.0

In most cases the crop areas selected using the DELP formulation are within one per cent of the DEMP results. In all but one case they are within two per cent. Total utility is also very accurately measured by DELP. The errors introduced by linearly segmenting the utility function appear negligible relative to the likely magnitudes of errors in the estimation of technical coefficients, utility function form and activity net revenues.

It is worth discussing the rationale for attempting to formulate LP versions of non-LP problems. A number of such formulations have been developed. For example, McCarl and Tice (1982) review a number of methods for approximating quadratic programming problems: MOTAD (Hazell 1971), separable programming (Thomas et al. 1972), marginal risk constraint LP (Chen and Baker 1974) and different forms of grid linearisation (Duloy and Norton 1975, McCarl and Tice 1980). They present three criteria by which to judge an approximation: the degree of error introduced, computational efficiency and human efficiency. It has already been argued that introduced error is not a significant problem with DELP. With regard to computational efficiency, there are positive and negative aspects to DELP. The positive aspect is use of LP rather than non-LP algorithms to solve the problem. On the other hand, DELP requires a larger matrix. Using McCarl and Tice's (1982) symbols (M = rows, N = columns, L = grid points, P = observations), DELP requires M + 2P rows and N + PL columns, compared with M + P and N + P for MOTAD. If L = 5 is accepted as adequate and 10 observations are used, the increase over a riskless LP model is only 20 rows and 50 columns.

DELP does require the calculation of a utility level for the wealth level corresponding to each utility activity. However, given the ease of constructing an electronic spreadsheet to conduct this task (or even to generate the model; see Pannell and Falconer 1987), the impact of this requirement on human efficiency need not be great.

Concluding Comments

Several points raised by Lambert and McCarl (1985) in their discussion of DEMP also apply to DELP.

(a) Any concave utility function can be used. It may represent increasing, constant or decreasing risk aversion.

(b) No assumptions about the underlying joint distribution of net revenues are required.

(c) A DELP model could be formulated to include variability in constraint limits (Lambert and McCarl 1985) or technical coefficients (Wicks and Guise 1978) or it could be included in a discrete stochastic programming model (Cocks 1968, Rae 1971).

(d) DELP requires the specification of a utility function. It does not deal with a class of functions, as does a stochastic dominance approach, except through repeated solutions.

DELP has a number of advantages relative to MOTAD, the most commonly cited method for risk programming with LP.

(a) It is consistent with expected utility maximisation, the basis of most economic risk theory.

(b) It accommodates any concave functional form for utility. Increasing, constant or decreasing risk aversion can be modelled.

(c) It requires no assumptions about the distribution of activity net revenues.

(d) Observed activity net revenues are entered directly into the matrix, so calculation and subtraction of mean revenues are not required. (However, see above comment on spreadsheets).

(e) Optimisation is achieved in a single step (or possibly two if grid points are re-focused). There is no need to derive an E-V frontier, requiring multiple solutions.

On the negative side, DELP does require a larger matrix than MOTAD and it requires the calculation of several utility levels for various wealth levels.

On balance, DELP offers several advantages over MOTAD at relatively minor cost. The method has been shown to allow use of LP to solve a non-LP problem with minimal introduced error.

References

Anderson, J.R., Dillon, J.L. and Hardaker, J.B. (1977), Agricultural Decision Analysis, Iowa State University Press, Ames.

Arrow, K. (1971), Essays in the Theory of Risk Bearing, Markham Publishing Co., Chicago.

Beneke, R.R. and Winterboer, R. (1973), Linear Programming Applications to Agriculture, Iowa State University Press, Ames.

Boussard, J.M. (1971), "A model of the behaviour of farmers and it application to agricultural policies", European Economic Review 2(0), 436-61.

Chen, J.T. and Baker, C.B. (1974), "Marginal risk constraint linear program for activity analysis". American Journal of Agricultural Economics 56(3), 622-27.

Cocks, K. (1968), "Discrete stochastic programming". Management Science 15(1), 72-79

Duloy, J.H. and Norton, R.D. (1975), "Prices and incomes in linear programming models", American Journal of Agricultural Economics, 57(4), 591-600.

Hazell, P.B.R. (1971), "A linear alternative to quadratic and semivariance programming for farm planning under uncertainty", American Journal of Agricultural Economics, 53(1), 53-62.

Hanoch, G. and Levy, H. (1969), "The efficiency analysis of choices involving risk". Review of Economic Studies, 36(3), 335-46.

Hey, J.D. (1979), Uncertainty in Microeconomics, Martin Robertson, Oxford.

Hey, J.D. (1981), Economics in Disequilibrium, Martin Robertson, Oxford.

Lambert, D.A. and McCarl, B.A. (1985), "Risk modeling using direct solution of nonlinear approximations of the utility function", American Journal of Agricultural Economics, 67(4), 846-52.

McCarl, B. and Tice, T. (1980), "Linearising quadratic programs through matrix diagonalisation", American Journal of Agricultural Economics, 62(3), 571-73.

McCarl, B.A. and Tice, T. (1982), "Should quadratic programming problems be approximated?", American Journal of Agricultural Economics, 64(3), 585-89.

McInerney, J.P. (1967), "Maximin programming - an approach to farm planning under uncertainty", Journal of Agricultural Economics, 18(2), 279-90.

Pannell, D.J. and Falconer, D.A. (1987), "Solution, interpretation and revision of MIDAS", in MIDAS: A Bioeconomic Model of a Dryland Farm System, ed. R.S. Kingwell and D.J. Pannell. Pudoc, Wageningen.

Rae, A.N. (1971), "An empirical application and evaluation of discrete stochastic programming in farm management", American Journal of Agricultural Economics, 53(4), 625-38.

Thomas, W., Blakeslee, L., Rogers, L. and Whittlesey, N. (1972, "Separable programming for considering risk in farm planning", American Journal of Agricultural Economics, 54(2), 260-66.

Wicks, J.A. Guise, J.W.B. (1978), "An alternative solution to linear programming problems with stochastic input-output coefficients", Australian Journal of Agricultural Economics, 22(1), 22-40.