Agricultural Economics, 6 (1991) 115–128
Elsevier Science Publishers B.V., Amsterdam

A stochastic dynamic programming framework for weed control decision making: an application to *Avena fatua* L.

Sushil Pandey [1]
Department of Agricultural Economics, School of Agriculture, The University of Western Australia, Nedlands, W.A. 6009, Australia

and R.W. Medd
Agricultural Research and Veterinary Centre, Forest Road, Orange, N.S.W. 2800, Australia
(Accepted 22 November 1990)

ABSTRACT

Pandey, S. and Medd, R.W., 1991. A stochastic dynamic programming framework for weed control decision making: an application to *Avena fatua* L. *Agric. Econ.*, 6: 115–128.

This paper develops a stochastic multi-period decision model to analyse a continuous wheat cropping system infested by wild oats (*Avena fatua* L.), in southern Australia. The multi-period solutions is obtained by employing a dynamic programming model in conjunction with a bioeconomic simulation model. An empirically estimated dose response function is used to derive the optimal herbicide rate. Uncertainties due to environmental effects on the performance of herbicide and crop yields are modelled and optimal decision rules derived. The results indicate that substantial economic gains can be realised if herbicide dose decisions are taken by considering future profit effects of current decisions, as opposed to the more common approach of only considering the current-period effect.

INTRODUCTION

Weeds impose a considerable economic burden on the Australian farmer. Total losses due to weeds for 1985/86 have been estimated to be approximately A$1500 million (Combellack, 1987). In 1984, farmers in Australia

[1] Present/Correspondence Address: Department of Agricultural Economics and Business Management, University of New England, Armidale, N.S.W. 2351, Australia.

spent over \$220 million on herbicides (Blacklow et al., 1984), a large proportion of which was for weed control in wheat. Despite such massive expenditure, may important attributes of the crop-weed-herbicide system are not explicitly modelled in deriving herbicide dose recommendations. For example, in some states of Australia it is illegal to apply any dose of herbicide other than that specified on the label (Pannell, 1989). In other states farmers may legally apply any dose up to that specified on the label but the recommendations are based on simplified singly-period threshold concepts with ad-hoc adjustments for risks.

Uncertainty in the crop-weed-herbicide system arises mainly from the variability in the performance of control measures, variability in the weed-free yield and variability of weed density. Feder (1979) has theoretically investigated the effect of some of these risks on the optimum dose of the control agent in a generic framework. Most contributions on modelling the effects of risk on pest control decisions are deficient because only one source of risk is modelled at a time, In reality, risks faced by farmers are manifestations of several sources of uncertainties and hence models which accommodate multiple sources of risk are likely to have greater relevance.

Benefits from weed control have a multi-period dimension due to (a) the effects of current level of control on future infestation levels, (b) development of herbicide resistance, and (c) herbicide carryover effects. Even though the current profits may not be adequate to recoup the current costs of treatment, some treatment may be justifiable if the possible prevention of future losses is also taken into account. Crop rotation decisions in many cases are governed by such long-term considerations. Similarly, if resistance to herbicides is likely to develop, recommendations based only on current period effects will be suboptimal. The models of decision making in weed control must thus be cast in a multi-period framework.

The purpose of this paper is to present a systems model for herbicide recommendations taking into account multi-period effects of current weed control decisions, stochastic influences and farmers' attitude towards risk. Our objective was to develop a stochastic multi-period decision model for weed control which could be solved using dynamic programming in conjunction with a bioeconomic model and which considered uncertainties in the performance of herbicide and crop yields in deriving optimal decision rules. An empirically estimated dose response function is used to derive the optimal dose of the control agent, a non-residual post-emergence herbicide. Although all three types of multi-period effects can be built into the model, data limitations permitted the incorporation of only the effect of the current level of control on future infestation. This is done in the context of a continuous wheat cropping system in southern Australia which is infested with wild oats (*Avena fatua* L.).

AN ECONOMIC MODEL OF WEED MANAGEMENT

For the purpose of this study, we have assumed a farm of non-specific size is uniformly infested with wild oats, is a closed system with negligible export or import of weed seeds and is managed by one farmer. The economic output of the farm is the grain yield of an annual crop. The farmer wishes to maximise profit (or utility) over a planning horizon of T periods. The profit function is assumed to be stationary in the sense that its parameters are constant over time. The decision problem is to derive an optimal strategy for the control of wild oats using diclofop-methyl. The solution is the dose of herbicide applied in each time period such that the present value over the planning horizon is maximised.

The solution to the problem can be found by using the tools of the optimal control theory. Let $B(\mathrm{SD}_t, X_t)$ define profit in time period 't' as a function of the seed density in the soil (SD) and herbicide dose (X). In the parlance of the optimal control theory, SD and X are the state and decision variables, respectively. The change in seed density from one time period to the next depends on the initial seed density which determines the potential for seed production during the current time period and the quantity of herbicide used which determines the reproductive output. Let $G(\mathrm{SD}_t, X_t)$ be a function measuring the change in seed density. The function G represents the equation of motion. Also, let $S(\mathrm{SD}_T)$ represent the terminal value of the weed seed bank at the end of the planning horizon. The objective function, assuming profit maximisation, is to maximise present value (PV):

$$\max \mathrm{PV} = \sum_{t=1}^{T-1} B(\mathrm{SD}_t, X_t)\delta_t + \delta_t S(\mathrm{SD}_T) \qquad (1)$$

subject to

$$\mathrm{SD}_{t+1} - \mathrm{SD}_t = G(\mathrm{SD}_t, X_t)$$

where δ_t is the discount factor for time period 't'. Applying Pontryagins' Principle of Maximum, one of the first-order conditions requires that:

$$\delta_t\, \partial B/\partial X_t + \lambda_{t+1}\, \partial G/\partial X_t = 0 \qquad t = 0, 1, \ldots, T-1 \qquad (2)$$

Equation (2) is equivalent to the usual first-order condition of setting marginal profits equal to zero if either λ_{t+1} or $\partial G/\partial X_t$ is assumed to be zero. The costate variable λ_{t+1} represents the marginal change in present value caused by a marginal change in the number of seeds at the beginning of time period 't'. Other things remaining constant, an increase in the current number of seeds (and hence weeds) reduces future profits; hence $\lambda_{t+1} < 0$. Also, $\partial G/\partial X_t < 0$ because, ceteris paribus, future weed population is reduced if the addition to the seed bank is reduced by increasing the current herbicide dose. The second term in equation (2), which measures the marginal benefit resulting from the effects of current level of control on future infestation, is hence non-negative. Due to the beneficial of the current level of control on future profits, the current marginal productivity of control inputs is increased. This results in a higher level of control than when current profits are maximised.

DYNAMIC PROGRAMMING MODEL

Dynamic programming is a computationally efficient method for solving the maximisation problem specified in equation (1). The method has been used for deriving optimal weed control strategies by Fisher and Lee (1981), Shoemaker (1982) and Taylor and Burt (1984). Its more general application for agricultural resource management has been reviewed by Kennedy (1986). The advantage of dynamic programming is that risk elements can be more easily incorporated, compared with other programming methods, and globally optimal solutions can be found even if the objective function is non-concave and discontinuous. To obtain a solution the total planning horizon is divided into periods (or 'stages' in the dynamic programming parlance) and the optimal solution for each stage derived. The interdependence of decision between stages is captured by using the concept of the state as these variables portray the nature of the system at the beginning of each stage. Thus the effects of decisions in one stage on the following stage is transmitted through the state variable. State variables need to be defined so that all information relevant to the current decision problem is embraced by the state variables. This requirement of dynamic programming is called the condition of Markovian independence (Nemhauser, 1966).

With the post-emergence herbicide as the control agent, weed density at the time of spraying is one of the state variables. If seeds exhibit dormancy, as in the case of wild oats, the number of viable seeds in the soil is another state variable. However, weed density can be ignored if a constant proportion of the seed bank is recruited. Thus, in the deterministic model in which all stochastic variables are replaced by their mean values, the seed bank population is the only state variable. The stochastic model is properly specified as having two state variables. However we used Taylor and Burt's (1984) decomposition method since it is more efficient for solving the stochastic model. The decomposition procedure is explained in a later section.

The uncertain variables included in the model are the efficacy of herbicide and weed-free yield. All other variables such as crop price, weed density and spray efficiency are assumed to be known with certainty. It has been the experience of farmers and researchers that the variability in herbicide efficacy is one of the dominant sources of risk in weed control. Similarly, weed-free yield is strongly influenced by climatic factors and this impinges on the profitability of weed control decisions.

In the literature on pest management, it is usual to consider pest density as an uncertain variable. Uncertainty arises both from non-uniform spatial distribution and with difficulties in measuring pest density. When the planning horizon is a single-period, it is straight forward to model spatial variability in weed density. However, in a multi-period framework, it becomes necessary to introduce a state variable for each portion of land with different weed densities. Such a dynamic programming model would be very difficult to solve. Thus the assumption made here is that a farmer would subjectively partition land into nominal categories according to the density of the weed. The spraying decision for a particular portion then depends on the weed density category. Formulated this way, the optimal decision rule for a farmer with more than one partition can be easily obtained from a single-state dynamic programming solution.

The solution procedure involved in the dynamic programming model is described by the recursive equation:

$$V_t(\mathrm{SD}_t, W_t) = \max_{X_t} \left[E\,\pi(\mathrm{SD}_t, W_t, X_t) + \delta\, \mathrm{E}\, V_{t-1}(\mathrm{SD}_{t-1}, W_{t-1}) \right]$$

$$t = 1, 2, \ldots, T \qquad (3)$$

where $V_t(\mathrm{SD}_t, W_t)$ is the optimal value function at stage 't' given seed density (SD) and weed density (W); π is the current profit if decision X is implemented and E is the expectations operator. In accordance with the dynamic programming method, time subscripts are specific in reverse order. Thus, the last year of the planning horizon is labelled as stage 1, the second last year as stage 2, and so on.

The length of the planning horizon, T, may be finite or infinite. In the solution of the infinite horizon problem, the optimal decision rule depends on the value of the state variable but not on the decision stages. In the case of the finite horizon problem the optimal decision depends both on the value of the state variable and the decision state. An approximate solution of a finite horizon problem can be derived by first solving the model for an infinite horizon such that $V_t \approx V_{t-1}$ and using the optimal decision rule corresponding to V_t for deriving solutions for a finite horizon problem.

In the deterministic model, weed density was deleted as a state variable because it is assumed to be a constant proportion of seed density. In the stochastic model weed density was retained along with seed density giving two state variables. The model was solved in two steps. First, the optimal value function for an infinite horizon problem was derived by dropping weed density as in the case of a deterministic model. In the second step, the optimal value function derived in the first step was substituted for V_{t-1} and an additional iteration solved. The recursive equation for the second step being:

$$V_t(\mathrm{SD}_t, W_t) = \max_{X_t} [E\,\pi(\mathrm{SD}_t, W_t, X_t) + \delta\, \mathrm{E}\, V_{t-1}(\mathrm{SD}_{t-1})] \tag{4}$$

The two-state variable problem is solved in the second stage by specifying current profit as a function of weed density and seed density. Although V_{t-1} is specified as a function of SD_{t-1} only, the decision rule derived is not myopic because the multi-period effects of current decisions are reflected by the seed bank population which appears in V_{t-1}. It is also assumed that weed density in the current period does not have any significant indirect effect on weed density in the next period.

For deriving numerical solutions, the state and decision variables were represented by 63 and 17 discrete values, respectively. The decision alternatives considered are different doses (including non-use) of diclofop-methyl up to the maximum permitted dose. For each starting value of the state variable, profits and ending values of the state variable were calculated for all discrete decision alternatives. For the ending value of the state variable falling between the two grid points, the optimal value function was approximated by linear interpolation between the adjacent grid points.

Risk is introduced to the model through random variables. Two random variables are required to incorporate risk of herbicide efficacy and weed-free yields. In the estimated weed kill function (see bioeconomic model), herbicide efficacy depends on soil moisture, which can take one of three ranked values. A discrete probability distribution for the soil moisture was derived by analysing climatic data. Values from this distribution were selected by Monte-Carlo sampling. In the case of weed-free yield, the simulated values (see bioeconomic model), were used directly as a sample from its distribution. The random variables associated with the soil moisture at the time of spraying and the weed-free yield were sampled independently. This is justified because correlation between them is expected to be low. Whereas herbicide efficacy is highly influenced by soil moisture at the time of spraying, the weed-free yield depends on the moisture regime throughout the life of the crop.

Bioeconomic simulation model

A bioeconomic simulation model was developed to trace the effect of weed control decisions on both the current and the future profits. The overall model is comprised of submodels for weed population dynamics, yield response to weed infestation, weed kill function, weed-free yield of crop, and climatic and economic factors.

A life cycle model of wild oats is used for predicting seedling recruitment, plant survival, seed production and seed survival. The recruitment of wild oats is not synchronised and occurs in waves during its life cycle (Quail and Carter, 1968; Amor, 1985). Given that the control agent is a non-residual post-emergence herbicide, modelling is helped by dividing seedlings into three cohorts. Plants which emerge before sowing belong to the first cohort. The second cohort encompasses plants emerging after sowing but before the post-emergence herbicide is applied. Plants emerging after the application of the post-emergence herbicide constitute the third cohort.

The seed bank is assumed to be homogeneous and recruitment in each cohort is specified as a constant proportion of the seed bank. Values of 23%, 12% and 2% recruitment were assumed respectively for the three cohorts. All seedlings in the first cohort are assumed to be killed by presowing operations. Due to the competitive effects exerted by seedlings upon each other, only a proportion of seedlings survive to maturity. Empirical evidence indicates that the majority of deaths in the second cohort seedlings occur before the biologically appropriate time for the application of the post-emergence herbicide (Medd, unpublished data). Thus, it is assumed that the full effect of density-dependent mortality is realised before the application of diclofop-methyl.

Plant fecundity is also density-dependent and is described by a hyperbolic function, where the maximum number of seeds produced per plant were 118, 22 and 4 for plants belonging to cohorts one to three, respectively. Some proportion of new seeds produced, in this case 10%, is assumed to be removed by the combine. Also, a proportion (65%) of the existing seed bank is lost due to natural mortality. Thus, the seed dynamics can be described by the following identity:

$$\mathrm{SD}_{t+1} = \mathrm{SD}_t - G_t - M_t + N_t \tag{5}$$

where SD_{t+1} is the size of the seed bank at the start of the period $t+1$, SD_t is the starting stock of seed bank, G is the loss due to recruitment, M is the loss due to mortality, and N is the new seed added to the seed bank. Most of the parameters for the model are obtained from experimental work at Orange, N.S.W. (Medd and Ridings, 1990). Values of the parameters unavailable from this source were obtained from experimental work in the United Kingdom (Cousens et al., 1986).

The yield (Y) of a weedy crop is specified as:

$$Y = Y^* g(W) \tag{6}$$

where Y^* is a parameter representing the maximum attainable yield in a weed-free situation given the level of environmental and management inputs, W is the weed density at maturity, and $g(\cdot)$ is a function often called the 'relative yield response' (Lanzer and Paris, 1981). By definition, $g(0) = 1$ and $g(\infty) = c^*$, where $0 < c^* \leq 1$. Thus the function $g(\cdot)$ provides a scaling factor. An implicit assumption contained in this specification is that Y^* and $g(\cdot)$ are separable. This assumption seems plausible and has been widely used to describe the yield response to various factors such as irrigation (Doorenbos and Kassam, 1979), fertilisers (Lanzer and Paris, 1981), weeds (Cousens et al., 1986) and insect pests (Feder, 1979).

It is usual to represent $g(\cdot)$ in the case of pests by a linear or sigmoidal function of pest density (Feder, 1979; Zimdahl, 1980). Cousens (1985) has argued, however, that $g(\cdot)$ is more accurately represented as a hyperbola in the case of weeds. The specific form suggested is:

$$g(W) = 1 - W/(a^{-1} + Wb^{-1}) \tag{7}$$

where 'a' and 'b' are the parameters, with values of 104.4 and 1.22, respectively. The parameter 'a' is a measure of the marginal yield loss as the weed density approaches zero. The parameter 'b' is an estimate of the maximum proportionate yield loss of a weedy crop. Since crops and weeds exert competitive effects on each other, yield loss due to weeds also depends on crop density. Based on Australian data, Martin et al. (1987) found 'a' to be proportional to crop density and their parameter estimates have been used herein.

Dose response function relates the quantity of herbicide applied to the proportion of weeds killed. The dose response relationship has the properties of a probability distribution function (Finney, 1971; Lichtenberg and Zilberman, 1986). Since the response to herbicides is a binomial variable (with the plants being considered as dead or alive), probit and logit regressions are the appropriate methods for efficient estimation of dose response relationships (Finney, 1971; Hewlett and Plackett, 1979). The logit specification is used in the present study, the specific relationship for diclofop-methyl being:

$$\log[P_i/(1 - P_i)] = \alpha_0 + \alpha_1 \log X_i + \alpha_2 \mathrm{SM}_i + \alpha_3 A_i + u_i \tag{8}$$

where P is the proportion of weeds killed, X is the quantity of herbicide applied (up to a maximum dose of 2 liters ha^{-1}), SM and A are measures of soil moisture and additives which also affect the performance of herbicide, and u is the random disturbance term.

Data from field experiments testing diclofop-methyl for control of wild oats, conducted by Hoechst in Western Australia, Victoria and New South Wales were adjusted for the effects of natural mortality using Abbott's formula (Finney, 1971). Based on the description of the soil moisture in the trial report as wet, average and dry, the soil moisture variable was rated as 3, 2 and 1, respectively. The dummy variable 'A' specified the addition ($A = 1$) or omission ($A = 0$) of wetting agent. Estimates of $\alpha_0 = -0.88 \pm 0.23$, $\alpha_1 = +1.66 \pm 0.24$, $\alpha_2 = +0.64 \pm 0.12$ and $\alpha_3 = +0.45 \pm 0.18$ ($\chi^2 = 151$, $n = 155$) obtained using the computer package GLIM (Anonymous, 1986) were all statistically significant ($P < 0.05$) and had the expected signs.

To take account of factors influencing the weed-free yield of wheat we employed a wheat growth simulation model developed at the Western Australian Department of Agriculture. The model inputs daily climatic data and yields for 74 years from 1912 to 1985 were predicted for Merredin, W.A. All management-specific inputs were assumed nonlimiting in the simulation. The yield predicted by the model was appropriately scaled down to reflect the average management practices. The average adjusted yield for the 74 simulated years was 1 t ha^{-1}.

RESULTS AND DISCUSSION

In order to derive a deterministic solution, the weed-free yield and the proportion of weeds killed for a given herbicide dose were set at their respective average values. As outlined under the dynamic programming methodology, the solution represents the optimal decision as a function of the state variable, but not of the decision state. Also, although the number of seeds in the soil is the state variable, results are presented in terms of weed density.

The optimal decision rule derived by solving the infinite horizon problem shows that herbicide dose increases, but at a diminishing rate, as weeds

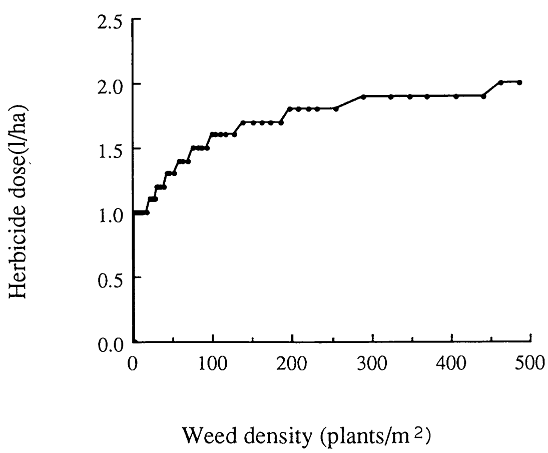

Fig. 1. Optimal deterministic multi-period decision rule.

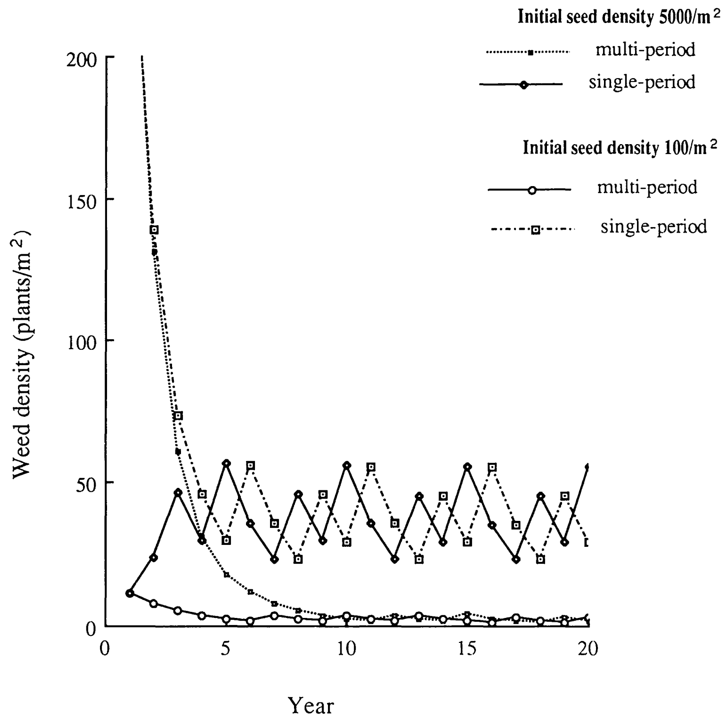

Fig. 2. Time traces of weed density resulting from the application of multi-period and single-period optimal decision rules for two initial seed densities.

become more dense (Fig. 1). The decision rule illustrated is straight forward and is seen as a useful applied decision making tool. A notable consequence of herbicide dose being dependent on weed density is that savings are possible. For instance, a herbicide dose of 1.5 litres ha^{-1} is specified for a wild oat density of 75 plants m^{-2} whereas 50 plants m^{-2} requires a dose of 1.3 litres ha^{-1}. Such savings affect profits and the present value of an infinite stream of profits is maximised if the decision rule is applied every year.

Time traces of weed density when the optimal decision rule (Fig. 1) is applied repeatedly are shown in Fig. 2. These traces were derived for initial seed densities of 100 and 5000 seeds m^{-2}, representing low and high populations. The weed density corresponding to the approximate steady-state in the case of multi-period optimisation was around 3 plants m^{-2} compared with 40 plants m^{-2} in the single-period case. Thus, by using the multi-period framework for decision making, the wild oat population is substantially minimised, but not eradicated at the optimal steady state.

The cumulative gain in present value when the multi-period, instead of the single-period decision rule is applied repeatedly, is shown in Fig. 3.

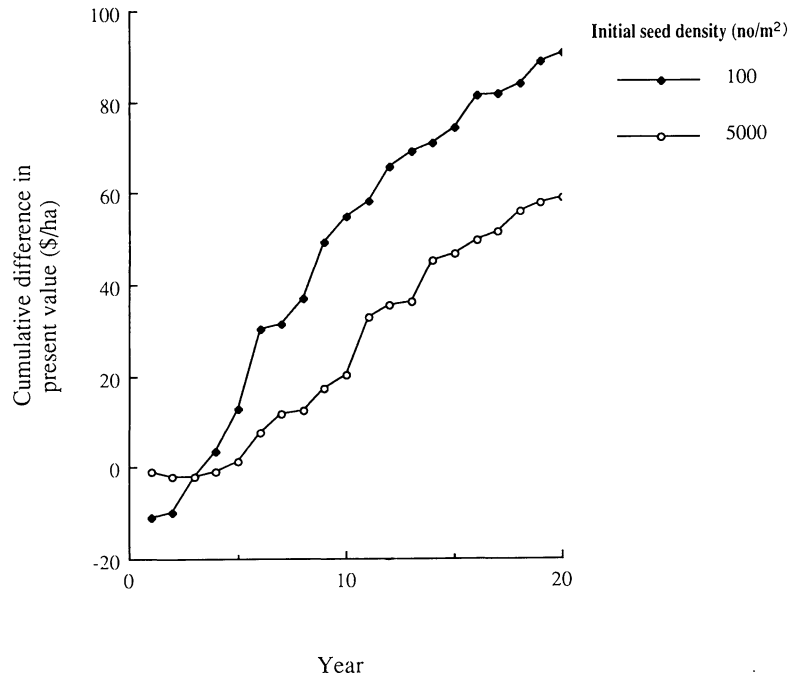

Fig. 3. Cumulative differences in present values between multi-period and single-period solutions for two initial seed densities.

These results demonstrate that the present value of profits is higher under multi-period optimisation and the gain over the single-period solution increases over time, more so at the lower initial seed density. Thus, in the case of wild oats, significant gains can be realised in future periods by reducing the weed burden early in the planning horizon even if current gains from such actions might be negative.

In the stochastic case, where transitions are probabilistic, solutions were derived using the two-step methodology described earlier and with the similar proviso that farmers were assumed to maximise expected profits. These results, along with the single-period optimal solution are depicted in Fig. 4. Here the optimal solution depends on weed density as well as seed density. For a given weed density, the optimal herbicide dose decreases with an increase in the seed bank. This is realistic since it is expected that the seed bank could be more readily manipulated by herbicide dose when there are few seeds in the soil to start with. On the other hand, when the initial seed bank is large, varying the addition to the seed bank by killing more weeds is unlikely to affect the seed bank substantially. When the seed bank is very large, the multi-period solution approaches that of the single-period.

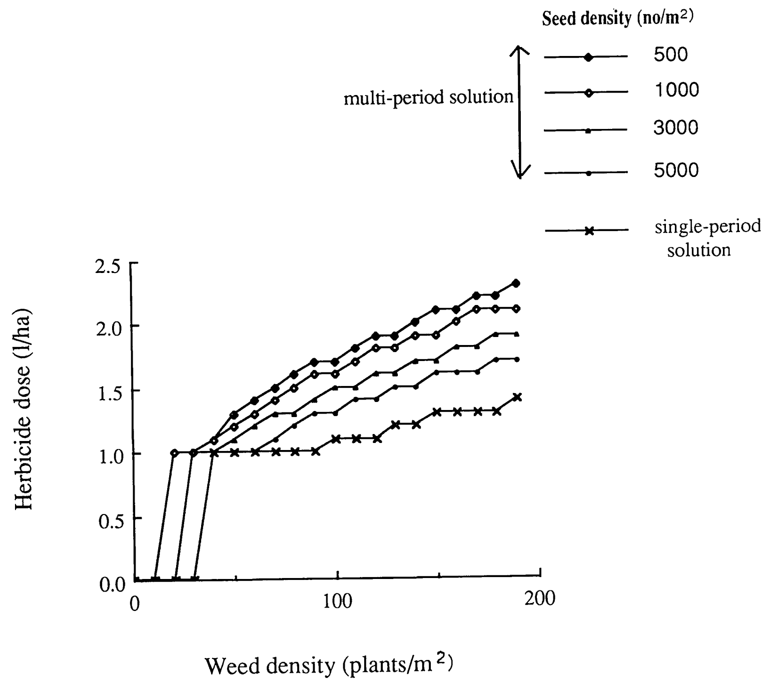

Fig. 4. Optimal stochastic multi-period and single-period decision rules.

If farmers are assumed to be risk-averse, then both the mean and variability of profits enter the objective function. If herbicide performance is assumed to be stochastic but the weed-free yield is deterministic, an increase in herbicide will reduce the variability of profits. Thus a risk-averse farmer would tend to apply more herbicide compared to a risk-neutral farmer. If herbicide performance is assumed to be deterministic but the weed-free yield is stochastic, the opposite result will hold because the variability of profits increases with an increase in herbicide rate. This is indicated by the structure of the yield response function specified in equation (6). The expected direction of change is ambiguous if both the weed-free yield and herbicide performance are assumed to be stochastic simultaneously. In this context, the generalisation usually made in the literature that herbicides are 'risk-reducing' (Binswanger, 1978; Feder, 1979) is therefore misleading. Our considerations reveal that the source of risk and how it is entered into the profit function influence whether or not an input such as herbicide reduces risk.

CONCLUSION

In order to improve decision making about herbicide dose in cropping systems we concluded it is necessary to consider multi-period effects and uncertainties due to environmental effects on crop yield, weed density and herbicide efficacy. This objective was achieved by combining a method of stochastic dynamic programming with a bioeconomic simulation model, to generate return matrix and transition probabilities, together with the novel inclusion of an empirically estimated dose response function. The development of such a model enabled the derivation of dynamically optimal dose of a post-emergence herbicide. The results indicated that the dynamically optimal solution maintained a lower steady-state weed population and higher economic returns compared with the single-period solution. From this we concluded that a long term approach to weed control is economically superior.

A limitation of the model is that only one of the long term effects, viz. The effect of current control on future infestation, is explicitly modelled. If resistance to herbicide is likely to develop, the policies derived here will not be optimal. However, the strength of the approach developed here is that such long term effects can be easily incorporated if the appropriate relationships can be quantified.

ACKNOWLEDGEMENTS

Funding from the Western Australian Wheat Research Committee and the University of Western Australia for Sushil Pandey, and from the Wheat Research Council of Australia for the demographic work of R.W. Medd conducted with the technical assistance of Kim Jones and Barry Kay at Orange is gratefully acknowledged.

REFERENCES

Amor, R.L., 1985. Seasonal emergence of weeds typically occuring in the Victorian cereal belt. Plant Prot. A., 1: 18–20.

Anonymous, 1986. The Generalised Linear Interactive Modelling Systems Manual. Numerical Algorithms Group, Oxford, 186 pp.

Binswanger, H.P., 1978. Risk and uncertainty in agricultural development: an overview. In: J.A. Roumasset, J.M. Boussard and I. Singh (Editors), Risk, Uncertainty and Agricultural Development. Agricultural Development Council, New York, pp. 383–397.

Blacklow, W.M., Pearce, G. and Humphries, A.W., 1984. Allocating scarce resources to weeds research in Australia. In: R.W. Madin (Editor), Proc. 7th Australian Weeds Conf., 17–21 September 1984, Perth, W.A., Australia, pp. 90–107.

Combellack, J.H., 1987. Weeds in cropping – their cost to the Australian economy. Plant Prot. Q., 24: 2.

Cousens, R., 1985. A simple model relating yield loss to weed density. Ann. Appl. Biol., 107: 239–252.

Cousens, R., Doyle, C.J., Wilson, B.J. and Cussans, G.W., 1986. Modelling the economics of controlling *Avena fatua* in winter wheat. Pestic. Sci., 17: 1–12.

Doorenbos, J. and Kassam, A.H., 1979. Yield response to water. FAO Irrig. Drain. Pap. 33, Food and Agriculture Organization, Rome, 193 pp.

Feder, G., 1979. Pesticides, information and pest management under uncertainty. Am. J. Agric. Econ., 61: 97–103.

Finney, D.J., 1971. Probit Analysis. Cambridge University Press, Cambridge, UK, 333 pp.

Fisher, B.S. and Lee, R.R., 1981. A dynamic programming approach to the economic control of weed and disease infestation in wheat. Rev. Market. Agric. Econ., 49: 175–187.

Hewlett, P.S. and Plackett, R.L., 1979. The Interpretation of Quantal Responses in Biology. Edward Arnold, London, 82 pp.

Kennedy, J.O.S., 1986. The econometrics of damage control: why specification matters. Am. J. Agric. Econ., 68: 261–273.

Lanzer, E.A. and Paris, Q., 1981. A new analytical framework for the fertilisation problem. Am. J. Agric. Econ., 63: 93–103.

Lichtenberg, E. and Zilberman, D., 1986. The econometrics of damage control: why specification matters. Am. J. Agric. Econ., 68: 261–273.

Martin, R.J., Cullis, B.R. and McNamara, D.W., 1987. Prediction of wheat yield loss due to competition by wild oats (*Avena* spp.). Aust. J. Agric. Res., 38: 487–499.

Medd, R.W. and Ridings, H.I., 1990. Feasibility of seed kill for control of annual grass weeds in crops. In: E.S. Delfosse (Editors), Proc. 7th Int. Symp. Biological Control of Weeds, 6–11 March 1988, Rome, Italy, pp. 645–650.

Nemhauser, G.L., 1966. Introduction to Dynamic Programming. Wiley, New York, 256 pp.

Pannell, D.J., 1989. Economics and the law relating to flexibility of chemical rates. Plant Prot. Q., 4: 104–106.

Quail, P.H. and Carter, O.G., 1968. Survival and seasonal germination of seeds of *Avena fatua* and *A. ludoviciana*. Aust. J. Agric. Res., 19: 721–729.

Shoemaker, C.A., 1982. Optimal integrated control of Univoltine pest population with age structure. Oper. Res., 30: 40–61.

Taylor, C.R. and Burt, O.R., 1984. Near optimal management strategies for controlling wild oats in spring wheat. Am. J. Agric. Econ., 66: 50–60.

Zimdahl, R.L., 1980. Weed crop competition: a review. International Plant Protection Centre, Oregon State University, 196 pp.