
Solving bioeconomic optimal control models numerically

O.J. Gacho

Department of Agricultural and Resource Economics, University of New England, Armidale

The influence of agricultural production decisions on the future state of the resource base, and thus on future productivity, demands the use of dynamic models for the analysis and design of efficient management strategies. Land use decisions affect the level of soil erosion, fertiliser and chemical runoff which may pollute waterways. Relatively simple simulation models can be used to explore the behavior of these production systems under alternative scenarios and management strategies. Simulation models can also be embedded into dynamic optimisation models to determine the levels of control variables over time which minimise or maximise a given objective function. The numerical solution of optimal control models, however, presents some problems which do not arise in plain simulation models, especially when nonlinear equations are present. This paper presents a model of a grazing system with three state variables (soil depth, soil fertility and pasture cover) and two controls (stocking rate and fertiliser application). The relative merits of solving the problem using either numerical optimal control, dynamic programming, or nonlinear programming are discussed.

Introduction

In recent years there has been a considerable increase in the number of dynamic optimisation studies reported in the agricultural and resource economics literature. Many of these models are bioeconomic in scope.

The biophysical simulation models which underlie these studies vary widely in their nature and complexity.

At one extreme are the simulation models developed by biologists and agricultural scientists, which tend to be very detailed in an attempt to make the model behavior approximate reality as closely as possible. Australian examples of these models include

GrassGro (Moore, Donnelly and Freer 1997) and APSIM (McCown et al. 1996); these models can be used to generate stochastic gross margin budgets, but their complexity precludes their use as part of optimising economic models.

At the opposite extreme, biophysical models developed by economists tend to be too simple by biologists' standards; these models are often linear or contain simple nonlinear functions to allow for the analytical solution of the optimisation problem and the derivation of comparative dynamic results — see, for example, Huffaker and Wilen (1991), LaFrance (1992), Goetz (1997) and Hu, Ready and Pagoulatos (1997).

Many of the variables included in complex simulation models are of no relevance from a decision making standpoint (that is, the decision rules are not sensitive to the values of these variables). Common-sense, experience and sensitivity analysis on large models can be used to identify those variables which can be excluded from the model and, eventually, a compromise may be reached between simplicity and realism which is mutually acceptable to economists and biologists.

Even with simple models the numerical solution of nonlinear optimal control problems is not trivial, and various simplifications (such as linearisation, nonlinear programming and dynamic programming) have been used; all these techniques require drastic model simplification (Kazmierczak 1996) and involve trade-offs. These issues are discussed and illustrated in this paper by developing and solving a model of a grazing system. According to Van der Ploeg et al. (1987), *bioeconomics* refers mainly to economic research using optimal control theory and models of population dynamics of exploited species. This definition is quite narrow compared with other interpretations of the term, but it does describe the work presented in this paper.

A grazing system provides a good example of a complex biophysical system because it exhibits interactions among plants, animals, the resource base (soil) and climate. The biophysical model of such a grazing system is developed here. It contains four state variables (soil depth, soil fertility, pasture mass and animal weight) and two control variables (stocking rate and fertiliser application rate), and is solved at a daily time step, to capture seasonal variations. The economic model, which is solved at annual intervals, attempts to identify the trajectories of the control variables through time which maximise discounted profits obtained from meat production. Meat is the only output considered for the sake of simplicity, but the model can be readily extended to account for wool production as well as other outputs, such as soil and fertiliser runoff. The model is simple enough to be described in a single paper yet detailed enough to provide a useful illustration of the essence of dynamic bioeconomic modeling.

The paper proceeds by defining an optimal control problem, describing the biophysical simulation model and presenting the derivation of a numerical production function. Three alternative solution techniques are then discussed and some results are presented. The paper ends with a few concluding comments regarding model limitations and future challenges.

The economic model

The economic model is based on a single-output production function and is an extension of the model presented by Cacho (1997). Meat is produced from grass and its production is indirectly controlled by stocking rate and fertiliser application. The amount of grass produced at any time depends on rainfall and the quality of the soil, described by soil fertility and soil depth. Soil fertility can be controlled directly through fertiliser application, while soil depth can only be controlled indirectly by maintaining a good pasture cover. Soil depth is determined by the balance between soil formation and soil loss, which in turn depend on pasture cover and rainfall. Because the model is dynamic, the effect of deteriorating soil quality on future productivity is automatically accounted for.

Table 1 defines the variables used in the economic model. The production function is assumed to be well

behaved (continuous, differentiable and concave in S and N — although N does not appear as an argument in the production function (equation 2), it affects meat production indirectly through its effect on fertility, which in turn affects pasture production (equations 4 and 5).)

The objective of the producer is to maximise discounted accumulated profits over the planning horizon. Ignoring the value of the land at the end of the planning horizon (T) the objective function is:

$$(1) \text{Max}_{S_t, N_t} J = \sum_{t=0}^T e^{-rt} [P_M M(t) - W_N N(t) - W_L L(t)] \Delta t$$

subject to:

$$(2) M_t = m(S_t, G_t)$$

$$(3) \Delta D_t = d(D_t, G_t, R_t)$$

$$(4) \Delta F_t = f(F_t, N_t, G_t)$$

$$(5) \Delta G_t = g(G_t, S_t, F_t, D_t, R_t)$$

$$(6) L_t = l(S_t, N_t)$$

$$(7) D_0 = D_K, F_0 = F_K, G_0 = G_K$$

$$(8) D_T \geq D_L, F_T \geq F_L$$

The term in square brackets in equation (1) is the annual profit obtained from meat sales net of fertiliser and labor costs. The maximisation of accumulated discounted profits is performed subject to the production function (2), the equations of motion (3) to (5), a labor requirement function (6), and the initial and final state constraints (7) and (8).

The equations of motion describe the dynamics of the biophysical system. The state variables soil depth (D_t), soil fertility (F_t) and pasture cover (D_t) interact among themselves and ultimately determine meat output. For any given rainfall pattern, the trajectory of the state variables through time is determined by the control variables stocking rate (S_t) and fertiliser application (N_t).

Notice that, although labor is normally treated as a primary input in economic models, the above optimisation is not performed with respect to this factor; rather, labor is treated as an intermediate input whose value depends on stocking rate and the amount of fertiliser applied, with both derivatives l_S and $l_N \geq 0$.

The Hamiltonian for this problem is:

$$(9) \quad H_t = e^{-rt} [P_M M_t - W_N N_t - W_L L_t] \\ + \lambda_{D_t} \Delta D_t + \lambda_{F_t} \Delta F_t + \lambda_{G_t} \Delta G_t$$

After dropping the time subscripts to avoid clutter, the solution is obtained by solving the system:

$$(10) \quad \frac{\partial H}{\partial S} = e^{-rt} [P_M m_S - W_L l_S] \\ + \lambda_{D_t} d_S + \lambda_{F_t} f_S + \lambda_{G_t} g_S = 0$$

$$(11) \quad \frac{\partial H}{\partial N} = e^{-rt} [P_M m_N - W_N - W_L l_N] \\ + \lambda_{D_t} d_N + \lambda_{F_t} f_N + \lambda_{G_t} g_N = 0$$

$$(12) \quad \frac{\partial H}{\partial D} = -\Delta \lambda$$

$$(13) \quad \frac{\partial H}{\partial F} = -\Delta \lambda$$

$$(14) \quad \frac{\partial H}{\partial G} = -\Delta \lambda$$

$$(15) \quad \frac{\partial H}{\partial \lambda_D} = d(D, G, R)$$

Table 1: Definition of variables

| | Variable | Definition | Units |
|--------------------------|-------------|-----------------------------------|-----------|
| Economic model | | | |
| Controls | S | stocking rate | hd/ha |
| | N | nutrient application | kg/ha |
| States | D | soil depth | m |
| | F | soil fertility | kg/ha |
| | G | pasture mass | kg DM/ha |
| Prices | P_M | price of meat | \$/kg |
| | W_N | price of fertiliser | \$/kg |
| | W_L | cost of labor | \$/h |
| Other | M | meat production | kg |
| | L | labor | h |
| | λ_j | costate variables | \$/yr/v a |
| Biophysical model | | | |
| | B | body weight | kg |
| | SF | soil formation | m/d |
| | SL | soil loss | m/d |
| | FG | fertility gain | kg/d |
| | FL | fertility loss | kg/d |
| | GG | pasture growth | kg/d |
| | GC | pasture cons. | kg/d |
| | R | rainfall | mm |
| | P_{cons} | pasture cons/hd | kg/d |
| | EI | energy intake | MJ/d |
| | EC | energy expenditure | MJ/d |
| | PI | proportion of max intake achieved | 1 |
| | E_{max} | maximum energy intake | MJ/d |
| | A | pasture allowance /hd | d |
| | I_{max} | maximum pasture intake | kg/hd |

a Units denoted by v represent either m (for λ_D) or kg/ha (for λ_F and λ_G).

$$(16) \quad \frac{\partial H}{\partial \lambda_F} = f(F, N, G)$$

$$(17) \quad \frac{\partial H}{\partial \lambda_G} = g(G, S, F, D, R)$$

$$(18) \quad \lambda_{DT} \geq 0, D_T \geq D_L, (D_T - D_L)\lambda_{DT} = 0$$

$$(19) \quad \lambda_{FT} \geq 0, F_T \geq F_L, (F_T - F_L)\lambda_{FT} = 0$$

$$(20) \quad \lambda_{GT} = 0$$

Note that constraint (8) forces the system to be sustainable with respect to D and F , by preventing the terminal value of these variables from falling below the lower limits D_L and F_L ; this constraint is translated into the Kuhn-Tucker conditions (18) and (19). Since the terminal value of G is not restricted, it is necessary to impose a constraint on the terminal value of its costate (λ_G) in order to obtain a unique solution, thus the need for equation (20).

The costate variables can be interpreted as dynamic shadow prices; each λ represents a measure of the future value of conserving the corresponding resource. The values of λ_j through time can provide valuable insight when plotted against their corresponding state variables in a phase diagram. This can be quite useful in policy analysis and in understanding the dynamics of the system.

The biophysical model

The biophysical model describes meat production and the dynamics of the state variables (3) to (5) as:

$$(21) \quad D_t = D_{t-1} + \Delta D_t$$

$$(22) \quad F_t = F_{t-1} + \Delta F_t$$

$$(23) \quad G_t = G_{t-1} + \Delta G_t$$

$$(24) \quad B_t = B_{t-1} + \Delta B_t$$

Note that B_t is a state variable in the biophysical model but not in the economic model, where this variable is part of the production function (2). Meat production forms a link between the economic and

biophysical models and is defined as:

$$(25) \quad M_t = B_t \cdot S_t$$

Although equations (21) to (24) are represented as single equations above, in the simulation model they are implemented as submodels, consisting of several equations each and, while the economic model is defined at an annual time step, the biophysical model is defined at a daily time step. This allows the model to capture seasonal variations in pasture production. Each submodel is described in turn in the following sections.

Soil quality

There are many characteristics of soil which determine its quality as a productive resource, including its composition, depth, structure and fertility. Here, soil quality is described by depth and fertility.

Soil depth

The variable D can be treated as a general measure of soil quality (that is, a productivity index), and could be interpreted as the most limiting soil characteristic other than fertility. However, the term *depth* is used for simplicity. It is assumed that at $D \geq 1$, soil quality is nonconstraining. The change in soil depth from one year to the next results from the cumulative balance between soil formation (SF) and soil loss (SL):

$$(26) \quad \Delta D_t = \int_{\tau=1}^{365} (SF_{\tau} - SL_{\tau}) d_{\tau}$$

where τ represents time measured in days. For convenience $\tau = 1$ can be taken to represent the beginning of the season and the season is assumed to last one year. The rate of soil formation (γ_D) is very slow and may be a function of soil depth (Goetz 1997, p. 343), but following McConnell (1983) here it is assumed to be constant, thus:

$$(27) \quad SF_{\tau} = \gamma_D$$

The rate of soil loss from a paddock depends on slope, pasture cover, rainfall and wind. For a given location and rainfall pattern SL can be described as a decreasing function of pasture cover:

$$(28) \quad SL_t = \alpha_D e^{-\beta_D G_t}$$

Soil fertility

The variable F can be regarded as a general measure of soil fertility, rather than as the concentration of a given nutrient (that is, the level of the most limiting nutrient). Soil fertility changes between years result from the balance between fertility gained (FG) and fertility used or lost (FL).

$$(29) \quad \Delta F_t = \int_{\tau=1}^{365} (FG_{\tau} - FL_{\tau}) d\tau$$

Fertility gain is controlled by nutrient application (N), but not all the fertiliser applied contributes to soil fertility, some may be lost to leaching or runoff. At low levels of pasture cover the proportion of fertiliser applied which contributes to fertility gain must be an increasing function of cover, because the presence of grass to capture nutrients and prevent water runoff will reduce fertiliser losses. The marginal reduction in fertiliser loss is expected to decrease as pasture cover increases, thus FG is assumed to increase at a decreasing rate as it approaches an asymptote:

$$(30) \quad FG_t = \left[K_F + \frac{\alpha_F G_{\tau}}{\beta_F + G_{\tau}} \right] N_t$$

Assuming that the amount of nutrients required to produce one unit of plant dry matter (γ_F) is constant, FL is represented as a linear function of pasture growth (GG):

$$(31) \quad FL_{\tau} = \gamma_F GG_{\tau}$$

Pasture mass

The change in pasture mass from one year to the next is the cumulative balance between net growth (GG) and consumption (GC) during the year:

$$(32) \quad \Delta G_t = \int_{\tau=1}^{365} (GG_{\tau} - GC_{\tau}) d\tau$$

Pasture growth at any time is a function of soil quality, environmental conditions and pasture mass:

$$(33) \quad GG_{\tau} = \alpha_G \frac{G_{\tau}^2}{Y_{\max}} \left[\frac{Y_{\max} - G_{\tau}}{G_{\tau}} \right]^{\gamma_G} . DE . FE . RE$$

Here, the first function on the right hand side represents net pasture growth on a given date under average conditions, while DE , FE and RE represent the effects of soil depth, soil fertility and rainfall on pasture growth. Seasonal variations in growth can be introduced by allowing the parameters α_G and Y_{\max} to exhibit cyclical annual patterns — see Cacho (1993) and Cacho, Bywater and Dillon (1998). Pasture growth is expected to increase at a decreasing rate with soil depth and fertility, thus:

$$(34) \quad FE = \frac{\alpha_{FG} F_{\tau}}{\beta_{FG} + F_{\tau}}$$

and

$$(35) \quad DE = \frac{\alpha_{DE} D_{\tau}}{\beta_{DE} + D_{\tau}}$$

The rainfall effect function can be used to introduce stochastic behavior into the model, but for a deterministic model it has a constant value and thus represents the effect of the average annual rainfall pattern:

$$(36) \quad RE = 1.$$

Pasture loss (consumption) is calculated as the average amount consumed per animal per day (P_{cons}) multiplied by the number of animals on the paddock:

$$(37) \quad GC_{\tau} = P_{\text{cons}} S_{\tau}$$

Other losses (such as senescence and decay) are embedded into equation (33) because GG represents net gain.

Animal growth

Animal growth results from the balance between energy intake (EI) and energy consumed (EC) for maintenance and other metabolic processes:

$$(38) \quad \Delta B_t = \int_{\tau=1}^{365} \psi_B (EI_{\tau} - EC_{\tau}) d\tau$$

where ψ_B is a conversion factor between energy and body weight. Energy intake is limited by the maximum amount of energy which the animal can con-

sume in a day (E_{max}) and by pasture availability:

$$(39) \quad EI_{\tau} = PI \cdot E_{max}, \quad 0 \leq PI_{\tau} \leq 1.$$

PI is a measure of pasture availability which takes account of competition between animals at high stocking rates, as described below. E_{max} depends on body weight and is described by the equation presented by Finlayson, Cacho and Bywater (1995):

$$(40) \quad E_{max} = \left[\alpha_B - \frac{\gamma_B B_{\tau}}{B_{max}} \right] B_{\tau}^{0.73}.$$

Pasture availability is expressed as the proportion of maximum intake actually achieved:

$$(41) \quad PI_{\tau} = \frac{A_{\tau}^2}{0.5 + A_{\tau}^2}$$

where A represents allowance, a measure of pasture available per animal relative to the requirements of the whole flock (or herd):

$$(42) \quad A_{\tau} = \frac{G_{\tau}}{I_{max} S_{\tau}}$$

and I_{max} is the intake limit per animal derived from energy intake and the amount of digestible energy available in the pasture (δ_G):

$$(43) \quad I_{max} = \frac{E_{max}}{\delta_G}.$$

The amount of energy consumed in metabolic processes is a function of body weight:

$$(44) \quad EC_{\tau} = \beta_B B_{\tau}^{0.75}.$$

The value of the exponent (0.75) is well accepted in the literature, whereas the parameter β_B depends on the status of the animal and on environmental conditions. This parameter can be a single value or a vector of values representing the annual cyclical pattern of energy requirements.

Finally, the amount of pasture consumed per animal (required by the pasture submodel) is estimated as:

$$(45) \quad P_{cons} = PI I_{max}.$$

The production function

The biophysical model represented by equations (21) to (45) must be solved by numerical integration and can be implemented in either a general purpose programming language (that is, Pascal, Fortran, C) or in a specialised simulation package. In this study the model was implemented and solved using Simulink® and Matlab® (Mathworks 1996a, 1997). The advantages of this software include: (i) the availability of a simple programming language for matrix manipulation; (ii) the ability to incorporate routines written in Fortran and C into the model; and (iii) the availability of a graphical environment to represent the biophysical model.

Figure 1: Simulated pasture mass and animal growth results for a 200-day season




Figure 2: Single-season production function produced by the biophysical simulation model

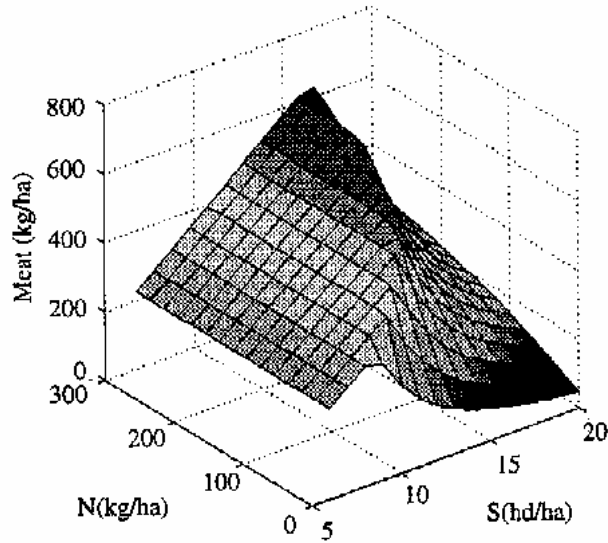


Table 2: Parameter values used in model runs

| | Parameter | Value |
|----------------|--------------------|---------------|
| Soil depth | γ_D | 0.0 |
| | α_D | 0.0002 |
| | β_D | 0.01 |
| Soil fertility | γ_F | 0.04 |
| | α_F | 1.0 |
| | β_F | 265 |
| | κ_F | 0.1 |
| | F_{max} | 2000 |
| | Pasture cover | α_G |
| γ_G | | 1.6 |
| Y_{max} | | 5000 |
| α_{FG} | | 1.0 |
| β_{FG} | | 50 |
| α_{DG} | | 1.0 |
| β_{DG} | | 0.5 |
| α_{RG} | | 1.0 |
| Weight | β_B | 0.8 |
| | α_B | 1.833 |
| | γ_B | 1.095 |
| | ψ_B | 0.0663 |
| | δ_G | 12.8 |
| | W_{max} | 55 |
| | Labor requirements | α_{LS} |
| α_{LN} | | 0.01 |
| Prices | P_M | 1.2 |
| | W_N | 0.1 |
| | W_L | 15.0 |

Using the parameters shown in table 2 and solving the biophysical model for a typical season lasting 200 days produces the results shown in figure 1. With an initial pasture mass of 1200 kilograms per hectare and an initial animal weight of 5 kilograms, pasture mass increases slightly and then decreases to 630 kilograms per hectare at the end of the season (figure 1A) as animals grow and consume grass to reach a final weight of 48 kilograms (figure 1B). In these simulations the control variables (S and N) are set at the start of the season. More complex models where the controls vary through the season are possible but are not considered here.

The single-season production function is estimated numerically by solving the biophysical model for a range of values of the control variables and using equation (25) to estimate meat output (figure 2). In effect, the simulation model is used to sample the production surface at the desired intervals. This process is analogous to performing a complete factorial experiment (Dillon and Anderson 1990, p. 68) with many levels of each factor. The production function illustrated in figure 2 forms the link between the economic and the biological models on the output side.

Up to this point the numerical solution of the model is trivial, relying simply on repeated numerical integration, which can be performed efficiently by a range of software packages. The solution of the

economic model, however, is considerably more complex, as discussed in the next section.

Solution techniques

Techniques used to solve optimal control models of the type described here fall within three broad categories: numerical optimal control (NOC), nonlinear programming (NLP) and dynamic programming (DP) and, within each category, there are various ways of implementing and solving the model. In this section, such techniques are reviewed and their advantages and disadvantages discussed.

Numerical optimal control

The NOC solution technique has appeal both because of its elegance and because it represents a direct translation of the mathematical model (10) to (20). A useful byproduct of this solution technique is the derivation of shadow price trajectories. In theory, the model can be solved numerically by an algorithm based on figure 3. In this compressed representation λ_t is the vector of costate variables [$\lambda_D, \lambda_F, \lambda_G$] and X_t is the vector of state variables [D, F, G] in year t .

Essentially, the direct NOC algorithm consists of three nested iterations with two search routines (figure 3). The algorithm starts by assigning initial values to the state and costate variables; however, the initial costate vector is not known, and an arbitrary value must be assigned to λ_0 to initiate the solution. In the inner loop, the biophysical model is run repeatedly for a single season, using different values

of S and N , until the combination which maximises H_t in equation (9) is found. There are several optimisation algorithms available to solve this part of the problem; Quasi-Newton techniques and the simplex method (Press et al. 1986) are two common examples. In practice, the first-order conditions (10) and (11) are not used in the algorithm because of difficulties in estimating numerical derivatives, as discussed later, and the presence of lower bounds in the controls, which can only take non-negative values.

Once the Hamiltonian has been maximised, the time counter is incremented, $\Delta\lambda$ is estimated via equations (12) to (14), and the costate vector (λ_t) is updated. H is then maximised for the new season and the iterative process continues until the end of the planning horizon (T) is reached; at this point the transversality conditions (18) to (20) are checked; if they are not satisfied the initial estimate of λ_0 is revised and a new iteration starts. This process continues until the value of λ_0 which satisfies all the transversality constraints is found. Search of the right λ_0 value for a given set of parameter values is time consuming and can be quite complex, especially when nonlinear interactions occur among the state variables, as is the case in this model.

The direct NOC solution is plagued with difficulties arising from the nonlinear and complex nature of the model. Estimation of the numerical derivatives in equations (12) to (14) is particularly troublesome. In principle, numerical differentiation is performed by taking small steps in state space (which consists of

Figure 3: Diagrammatic representation of the numerical optimal control algorithm

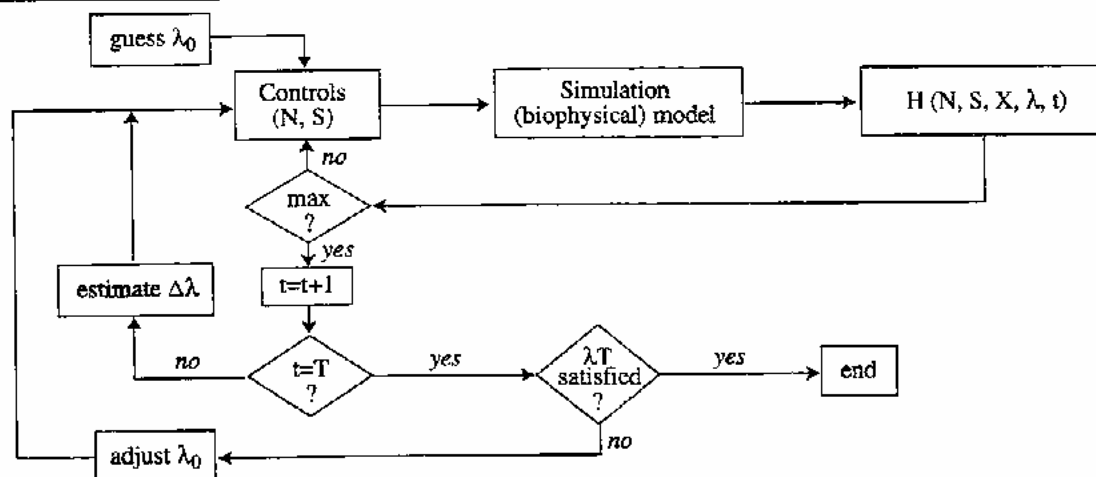
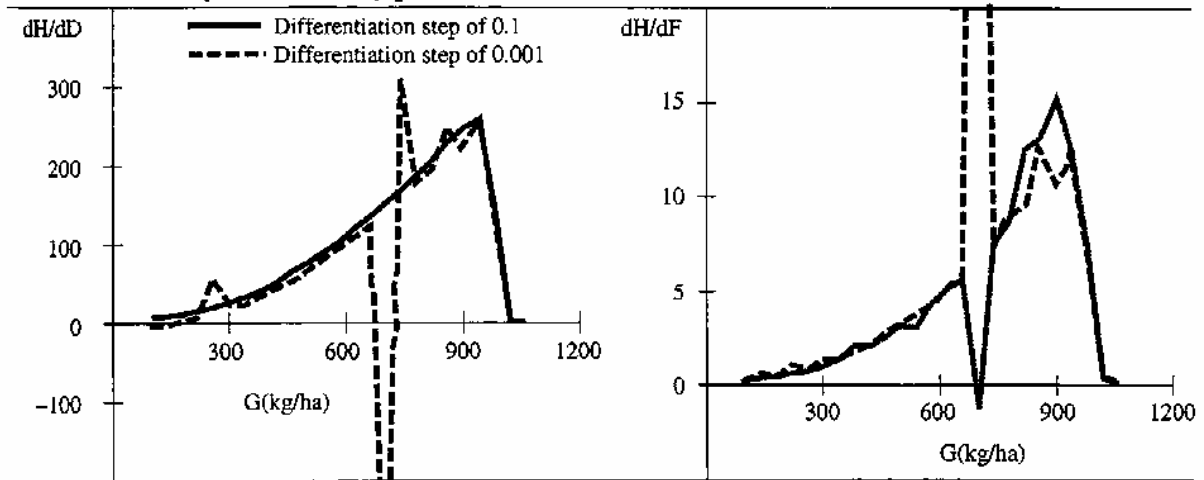


Figure 4: Estimated partial derivatives of the Hamiltonian with respect to soil depth and soil fertility as affected by pasture mass



solving the model for alternative values of the state variables and taking first differences). In practice, the results obtained from single differencing may be poor approximations of the derivatives and are highly sensitive to the step size (h) used. Thus, it is necessary to use techniques based on a Taylor series expansion (see, for example, Linfield and Penny 1995, pp. 120–4). With complex nonlinear models, even these more precise numerical differentiation techniques may encounter problems in certain areas of the production surface, as illustrated in figure 4.

An approximation using a third-degree expansion (based on first, second and third differences) produces an error of the order h^4 and can handle most situations, but not all. A troublesome region in the state surface exists at pasture masses in the neighborhood of 700 kilograms per hectare (figure 4). This appears to be a critical point, where small changes in stocking rate can make the difference between pasture persistence and disappearance. Thus the size of h must be carefully selected in regions such as this, or additional constraints on allowable state variable values must be introduced.

Recent research has focused on ‘control parametrisation’, which approximates the optimal control problem by a constrained NLP problem (Jennings et al. 1997), and other discretisation techniques (Schwartz 1996). Given the complexity of this literature, and the level of programming skills required, it is not surprising that most empirical bioeconomic models have been solved using simplifications based on DP or NLP models. Developers of specialised

optimal control software (Jennings et al.; Schwartz) claim that their packages can solve large and difficult problems. The two packages available have not been tested with the present model so these claims can neither be supported nor rejected.

Dynamic programming

The DP approach simplifies the problem by breaking it down into a series of small subproblems (Bellman and Dreyfus 1962; Dreyfus and Law 1977). For the type of problem discussed here the DP technique is based on the repeated solution of the recursive equation:

$$(46) V_t(\mathbf{X}_t) = \max [R_t(\mathbf{X}_t, \mathbf{U}_t) + \delta V_{t+1}(\mathbf{X}_{t+1})]$$

where \mathbf{X}_t is a vector of state variables, \mathbf{U}_t is a vector of control variables and δ is a discount coefficient. In our problem $\mathbf{X}_t = [D_t, F_t, G_t]$ and $\mathbf{U}_t = [S_t, N_t]$. V_t is also called the ‘optimal value function’ and R_t is the one-period return function, represented by the expression in brackets in equation (1).

The DP technique can solve unruly problems and can be used to derive estimates of the shadow prices, as explained by Kennedy (1988, p. 62). Another advantage of the DP algorithm is that it produces a decision rule which can be saved to disk and later used to quickly determine the optimal trajectory for different initial conditions. In contrast, other techniques require the model to be solved whenever the initial conditions change.

The main disadvantage of the DP approach is its vulnerability to the ‘curse of dimensionality’, which

refers to the explosive growth of these models as the number of states increases. The recursive equation is solved at discrete intervals in state space and backwards in time. Specifically, a 'state grid' is created by defining a multidimensional array of points, where each dimension represents a state variable. The biophysical model is repeatedly solved at each point on the grid in order to find an approximation to the optimal control (U_t^*) at the given state X_t . For nonlinear problems the accuracy of the result depends critically on the coarseness of the state grid.

To illustrate the dimensionality problem, consider sampling at five values for each state in our problem. Given the three state variables, the resulting grid will have 125 points per time period, which means that, for a 16-year planning horizon the recursive equation must be solved 2000 times. But note that solving the recursive equation (46) requires repeated solution of the biophysical model to find a maximum.

The simplest approach to the maximisation problem consists of sampling the control space at regular intervals and selecting the control vector with the highest return in the set. Given our two controls and sampling at six values of N and S , we have 36 function evaluations per state, which translates into a total of 72 000 (that is, $16 \times 5^3 \times 6^2$) runs of the annual biophysical model being required to solve the DP model. Such a coarse grid, however, is likely to result in a poor approximation to the optimal trajectory. Increasing the resolution of the model by sampling

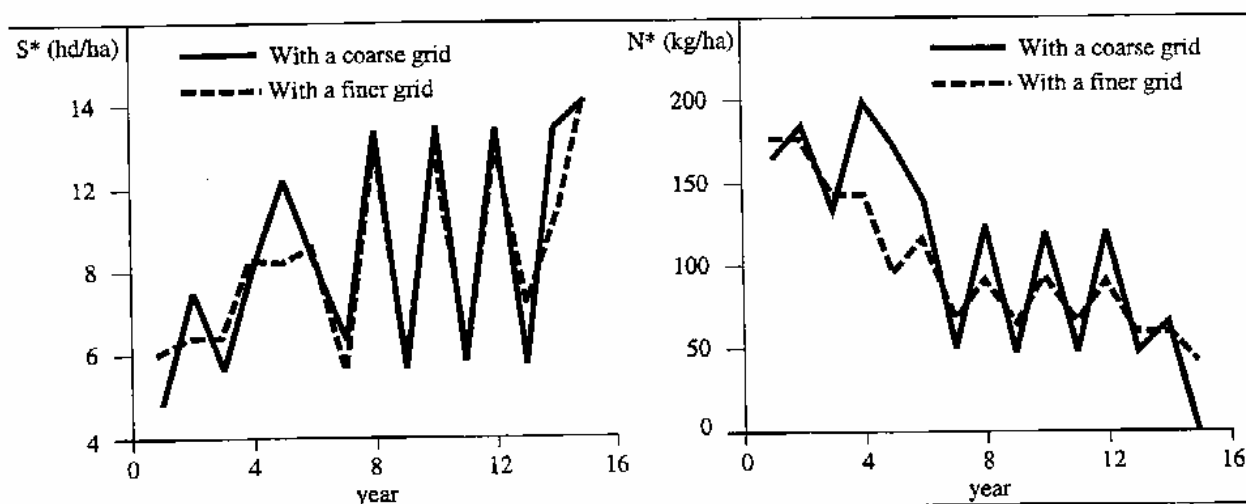
at 9 points per state and 10 points per control results in 1 166 400 (that is, $16 \times 9^3 \times 10^2$) runs of the biophysical model.

In the latter case, if it takes one minute to complete one run of the biophysical model, it would take 810 days to obtain a solution, and this would be for a single run of a deterministic model. Now consider the additional costs of running sensitivity analyses on prices, costs and discount rates, never mind attempting to make the model stochastic. Clearly, the biophysical model must be dramatically simplified to obtain a solution within an acceptable time span.

The accuracy of the DP model for a given grid size can be dramatically improved by using linear interpolation when evaluating equation (46). Note that, given the discrete nature of the sampling space, in evaluating the return function for a given value of U_t it is unlikely that the terminal state X_{t+1} will coincide with a point on the grid, which means that the required value of V_{t+1} may not be available. The easy solution is to use rounding or truncation and use the closest V_{t+1} value available, but with a coarse grid this may produce unacceptable results.

Linear interpolation between sample points is a numerically efficient alternative which can improve the accuracy of the model considerably. Another way of improving the accuracy and efficiency of the model is to use an optimisation algorithm, rather than uniform sampling, to find $V_t(X_t)$. Also, for some nonlinear problems, defining the state grid at unequal

Figure 5: Approximated optimal control trajectories obtained by solving the DP model



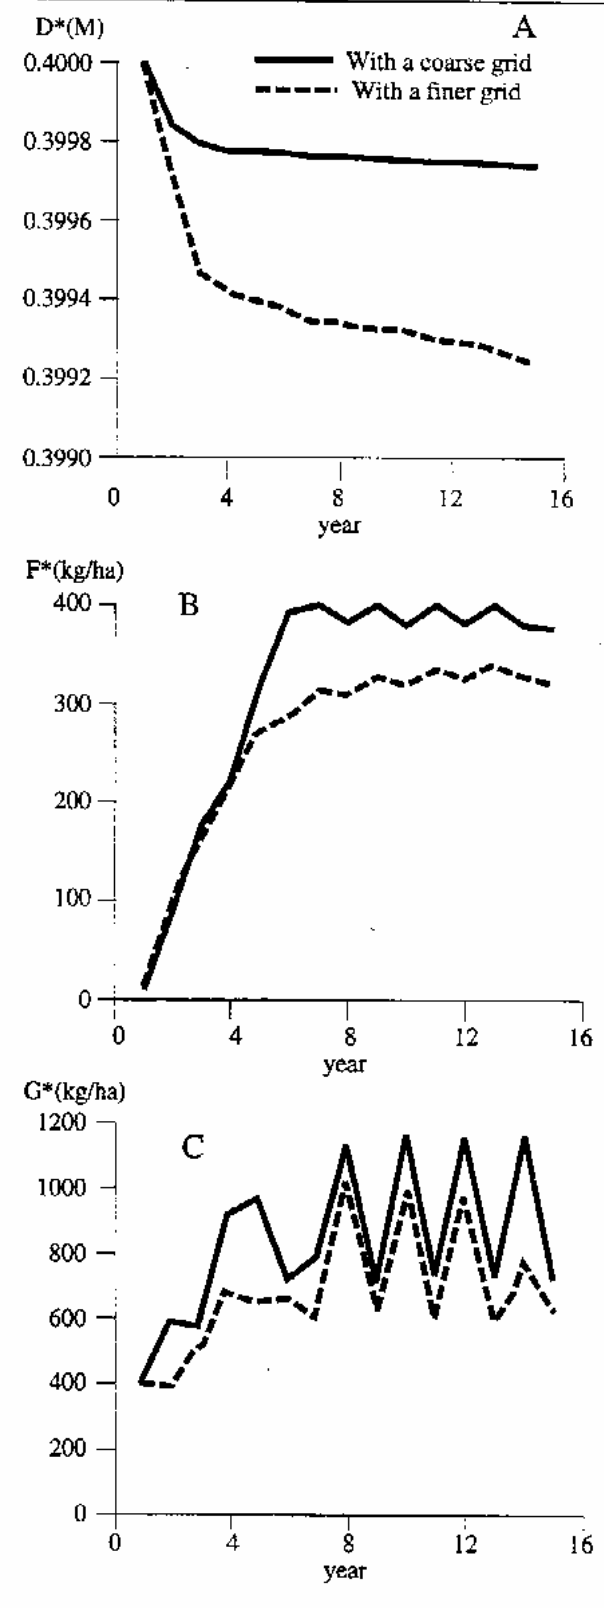
intervals or on a logarithmic (rather than linear) base may prove beneficial.

The DP problem was solved for a 16-year planning horizon using the parameter values in table 2. A relatively degraded paddock was simulated, represented by the initial state $(D_0, F_0, G_0) = (0.4, 10, 400)$. The same problem was solved using two different grid sizes: (5,5,5) and (9,9,9) and two corresponding sets of control sampling points (6,6) and (10,10). To improve accuracy, linear interpolation was used in estimating equation (46) and, to speed up the model solution, a multidimensional production function, estimated in advance and saved to disk, was used to approximate output. This amounts to using the (numerical) production function as a look-up table so that, instead of running the biophysical model repeatedly, the meat output corresponding to a given set of U and X values is found by interpolation. The equivalent technique for a stochastic model would involve estimating a multidimensional transition probability matrix (the stochastic equivalent of the numerical production function described above), a topic which is beyond the scope of this paper.

For the assumed initial conditions, the estimated optimal control trajectories (figure 5) consist of low stocking rates and high fertiliser rates early in the planning horizon (to improve the productivity of the soil), followed by higher stocking rates and lower fertiliser rates in later years. This control strategy results in decreasing soil depth and increasing soil fertility (figures 6A and 6B), with both state variables tending to stable long run equilibriums. Pasture cover appears to tend to an average of about 900 kilograms per hectare (figure 6C), with wide year-to-year variations caused by changes in S .

The jagged shape of the optimal trajectory is caused by the coarseness of the control grid, a smoother trajectory is obtained as the control grid becomes finer (see dashed line in figure 5), with a consequent cost in solution time. It is obvious from these results that the nature of the problem requires a finer sampling grid in control space, so it would be more efficient to use a search routine to solve for the optimal value function rather than sampling at uniform intervals in control space; the same type of quasi-Newton and simplex techniques discussed earlier can be used here.

Figure 6: Optimal trajectories of state variables obtained by solving the DP model



Non-linear programming

A common way of solving optimal control models is to redefine them as NLP models. In general terms, this technique consists of defining the state and control variables in each time period as activities, and specifying the equations of motion as nonlinear constraints linking variables across time periods (Standiford and Howitt 1992). The most common software tool used in published reports of this type appears to be GAMS/MINOS (Brooke, Kendrick and Meeraus 1988).

With this technique, the number of decision variables in the NLP model is determined by the number of control and state variables multiplied by the number of time periods in the planning horizon. Thus, solving our problem for 16 years we would have a minimum of 96 ($16 \times 3 \times 2$) activities and 48 (16×3) constraints, not counting other auxiliary activities and constraints required by the problem. Hertzler (this volume) points out that only inequality constraints need to be explicitly included in the NLP model since equality constraints can become part of the simulation. This reduces the dimensionality of the problem considerably (to 32 activities in our example), at the cost of not obtaining estimates of shadow prices.


With a software package as powerful as GAMS, NLP models are fairly easy to implement, but with complex nonlinear models they may fail to find a global maximum, as shown by Michalewicz (1992,

pp. 112–18). To illustrate the problem of finding the global maximum, the model was solved as an NLP problem using a sequential quadratic programming algorithm (Mathworks 1996b); the results are shown in figure 7. Initially it was attempted to solve the problem for 16 years (32 decision variables), but the model did not converge after a few thousand iterations. When the planning horizon was decreased to 10 years the solution was obtained in 56 minutes. It required approximately 800 iterations to reach convergence, this represents 8000 runs of the biophysical model, because each iteration of the NLP model represents a 10-year run.

The NLP model was solved for both a healthy system — $(D_0, F_0, G_0) = (1.0, 150, 1200)$ — and a fairly degraded system — $(D_0, F_0, G_0) = (0.4, 10, 400)$. The optimal control trajectories (S^* and N^*) obtained in the degraded system (figure 7) were similar to those obtained with the DP model, starting with a low S (4 head per hectare) and high F (110 kilograms per hectare), to allow the system to recover, and ending with high S (14 head per hectare) and a relatively low N (60 kilograms per hectare). In a healthy system, the optimal S trajectory remained high (around 14 head per hectare) and the optimal N trajectory decreased from an initial value of 200 kilograms per hectare to a final value of 30 kilograms per hectare.

When the final value of pasture cover (G_T) was constrained to be no less than 600 kilograms per hectare (to prevent a totally degraded system to be left at the end of the planning horizon), the NLP algo-

Figure 7: Optimal trajectories of the control variables obtained from running the NLP model



rithm had a tendency to get stuck in situations where this constraint was violated (with values around 50 kilograms per hectare).

Summary and conclusions

Due to space and time constraints, several simplifications were used and shortcuts taken in the exposition of the problem and solution techniques. The most obvious omission is the final value of the land, which is likely to be influenced by the quality of the soil. The fact that this price was neglected amounts to an assumption that the land market does not respond to soil quality. This omission can be easily remedied by adding a term to the end of equation (1). Also, the parameter values used were set arbitrarily, based on commonsense and previous experience of the author, to yield reasonable results. The system simulated here represents a generic situation using an animal with growth characteristics similar to a sheep grazing on a fairly productive grass. This is not seen as a limitation of the model, but as a means of collaborating with agricultural scientists who may produce the research required to estimate the biophysical parameters. Finally, the cost of changing stocking rate was ignored (that is, the need to purchase animals or raise replacements); this cost can be included in the profit function of equation (1).

Readers who may be bothered by the informal style of the paper are reminded that there are many formal expositions of optimal control problems in the literature, and the main purpose of this paper is to present practical information which is not available elsewhere. Mathematical jargon was avoided to the extent possible to make the paper accessible to a larger audience. Dynamic bioeconomic models have much to contribute in policy analysis and in the design of sustainable production systems, but we are just scratching the surface.

References

Bellman, R.E. and Dreyfus, S.E. 1962, *Applied Dynamic Programming*, Princeton University Press, New Jersey.

Brooke, A., Kendrick, D. and Meeraus, A. 1988, *GAMS a User's Guide*, Scientific Press, San Francisco.

Cacho, O.J. 1993, 'A practical equation for pasture growth under grazing', *Grass and Forage Science*, vol. 48, pp. 387–94.

— 1997, Dynamic models, externalities and sustainability in agriculture. Contributed paper, 41st Annual AARES Conference, 20–25 January, Gold Coast, Australia.

—, Bywater, A.C. and Dillon, J.L. 1998, 'Assessment of production risk in grazing models', *Agricultural Systems*.

Dillon, J.L. and Anderson, J.R. 1990, *The Analysis of Response in Crop and Livestock Production*, Pergamon Press, Oxford.

Dreyfus, S.E. and Law, A.M. 1977, *The Art and Theory of Dynamic Programming*, Academic Press, New York.

Finlayson, J.D., Cacho, O.J. and Bywater, A.C. 1995, 'A simulation model of grazing sheep: I. Animal growth and intake', *Agricultural Systems*, vol. 48, pp. 1–25.

Goetz, R.U. 1997, 'Diversification in agricultural production: a dynamic model of optimal cropping to manage soil erosion', *American Journal of Agricultural Economics*, vol. 79, pp. 341–56.

Hu, D., Ready, R. and Pagoulatos, A. 1997, 'Dynamic optimal management of wind-erosive rangelands', *American Journal of Agricultural Economics*, vol. 79, pp. 327–40.

Huffaker, R.G. and Wilen, J.E. 1991, 'Animal stocking under conditions of declining forage nutrients', *American Journal of Agricultural Economics*, vol. 73, pp. 1213–23.

Jennings, L.S., Teo, K.L., Fisher, M.E. and Goh, C.J. 1997, *MISER3 Optimal Control Software, Version 2.0, Theory and User Manual*, Department of Mathematics, University of Western Australia, Perth.

Kazmierczak Jr, R.F. 1996, *Optimizing Complex Bioeconomic Simulations Using an Efficient Search Heuristic*, DAE Research Report no. 704, Louisiana State University, Baton Rouge.

Kennedy, J.O.S. 1988, 'Principles of dynamic optimization in resource management', *Agricultural Economics*, vol. 2, pp. 57–72.

LaFrance, J.T. 1992, 'Do increased commodity prices lead to more or less soil degradation?' *Australian Journal of Agricultural Economics*, vol. 36, pp. 57–82.

Linfield, G. and Penny, J. 1995, *Numerical Methods Using Matlab*, Ellis Horwood, New York.

McConnell, K.E. 1983, 'An economic model of soil conservation', *American Journal of Agricultural Economics*, vol. 65, pp. 83–9.

-
- McCown, R.L., Hammer, G.L., Hargreaves, J.N.G., Holzworth, D.P. and Freebairn, O.M. 1996, 'APSIM: a novel software system for model development, model testing, and simulation in agricultural systems research', *Agricultural Systems*, vol. 50, pp. 255-71.
- Mathworks 1996a, *Using Matlab*, The Mathworks Inc., Natick, Massachusetts.
- 1996b, *Matlab Optimization Toolbox, Users Guide*, The Mathworks Inc., Natick, Massachusetts.
- 1997, *Using Simulink*, The Mathworks Inc., Natick, Massachusetts.
- Michalewicz, Z. 1992, *Genetic Algorithms + Data Structures = Evolution Programs*, Springer-Verlag, Berlin.
- Moore, A.D., Donnelly, J.R. and Freer, M. 1997, 'GRAZPLAN: decision support systems for Australian grazing enterprises. 3: Pasture growth and soil moisture submodels, and the GrassGro DSS', *Agricultural Systems*, vol. 55, pp. 535-82.
- Press, W.H., Flannery, B.P., Teukolsky, S.A. and Vetterling, W.T. 1986, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press, Cambridge.
- Schwartz, A. 1996, Runge-Kuta integration for solving optimal control problems, PhD thesis, University of California, Berkeley.
- Standiford, R.B. and Howitt, R.E. 1992, 'Solving empirical bioeconomic models: a rangeland management application', *American Journal of Agricultural Economics*, vol. 74, pp. 421-33.
- Van der Ploeg, S.W.F., Braat, L.C. and Van Lierop, W.F.J. 1987, 'Integration of resource economics and ecology', *Ecological Modelling*, vol. 38, pp. 171-90.

Gooday, J. (ed.) 1998, *Proceedings of the Bioeconomics Workshop*, Workshop (Post-Australian Agricultural and Resource Economics Society Conference) sponsored by ABARE, 22 January, Armidale, New South Wales, ABARE, Canberra.

Australian Bureau of Agricultural and Resource Economics
GPO Box 1563 Canberra 2601

Telephone +61 2 6272 2000 Facsimile +61 2 6272 2001
Internet www.abare.gov.au

ABARE is a professionally independent government economic research agency.

The papers in this proceedings were presented at the Bioeconomics Workshop, post the Australian Agricultural and Resource Economics Society Conference, 22 January 1998, held at the University of New England, Armidale, New South Wales, Australia. The papers in this proceedings have been peer reviewed **in their entirety** and revised following presentation at the conference.

The referees are gratefully acknowledged for their comments and suggestions.