



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

AN ADAPTIVE CONTROL APPROACH TO AGRICULTURAL POLICY*

J. W. FREEBAIRN and G. C. RAUSSER

*Australian National University and University of
California, Davis*

It is postulated that some issues of economic policy in general, and of Australian agricultural policy in particular, may be analysed in the framework of an adaptive control model. Policy making is characterized as a rational, sequential decision-making process under conditions of imperfect knowledge in which forthcoming information may be used to learn about the uncertain terms as decision periods pass. Emphasis is given to the linear-quadratic control problem. The paper provides a review of the formulation of a policy problem in the framework of an adaptive control model and of derived policy strategies. An illustrative example is reported.

Introduction

An adaptive control model analysis of economic policy formalizes the process of policy making as a multiperiod optimization problem under conditions of imperfect knowledge. In particular, the model recognizes uncertainty about the effects of alternative policy actions. As the process proceeds through decision periods additional information becomes available. The new information is used in each decision period to specify the policy action for that period and to learn more about the effects of alternative policy actions. Further, the adaptive control model recognizes that current period decisions can actively influence the information generated for learning. This paper reviews some recent literature embracing the applicability of adaptive control theory for the analysis of policy making and suggests how the models could be applied to the analysis of some problems of Australian agricultural policy.

The paper is structured as follows. The next section discusses the underlying structure and the results of an adaptive control model analysis of a policy problem. A specific class of control problems is considered in a formal way in the subsequent section. These problems are characterized by a quadratic objective function and linear constraints with probabilistic (as opposed to single valued or perfect) knowledge about the values of the parameters and some of the variables of the constraint functions. This model is employed in the analysis of an illustrative version of Australian wheat supply and price policy. A final section provides a summary and some concluding comments.

Economic Policy in an Adaptive Control Framework

The policy problem is formulated as one of choosing 'desirable' values for a set of policy or instrument variables. In the context of an adaptive control model the 'desirable' values are obtained from the solution of a formal optimization problem.¹ This section discusses: (1) the repre-

* Giannini Foundation Research Paper No. 384.

¹ This framework of policy analysis stems from the early work by Tinbergen [11].

sentation of a policy problem in the framework of an optimization problem, (2) the nature of the derived policy actions, and (3) the representation of Australian wheat stabilization policy in the framework of an adaptive control model.

Policy Problem

The formulation of a policy choice problem as an optimization problem involves four related components, namely, the policy variables, the objective function, the constraints describing physical evolution of the economic system over time, and the process of information collection and learning over time.

Specification of the policy variables enumerates those variables which may be manipulated by the policy maker(s) and also the procedures and points in time at which these variables may be revised in the light of forthcoming information. The form and extent of potential intervention in the economic system is made explicit.

The objective function of the model ranks the desirability of different economic states which are directly or indirectly influenced by alternative settings of the policy variables. Several argument variables may be required to describe performance of the economic state. Termed performance variables, these would be variables which are thought by the policy maker(s) to have significant welfare connotations for different members of society and would include policy variables and endogenous variables influenced by the policy variables. Given the diversity of preferences of different sectors of society it would seem reasonable to use a set of objective functions rather than a unique function in the analysis. In particular, the functions would span a range of trade-offs between various arguments of the function to reflect the relative intensity of preferences of different sectors of society.²

The constraint functions describe the policy possibility set. A set of initial conditions define the present state of the economic system. Evolution of the system over time is described by state transformation functions relating future period states to policy variables, other exogenous variables and current state variables. These functions provide an important dynamic component to the policy problem and provide a link by which current decisions influence the future as well as the current period performance of the system. The state transformation functions are based on a mathematical model of the main causal relationships describing behaviour of the system.

The model allows for uncertainty about the effects of alternative policy actions via imperfect knowledge about some of the parameters and variables of the state transformation functions. It is assumed that policy makers have probabilistic information about these uncertain terms, and further, that forthcoming information may be used to update these probability distribution functions.

An important component of the problem concerns the collection of new information as the process proceeds and the way in which this information is used to learn about the uncertain components of the constraint functions. The former involves enumeration of the new information collected, and the latter embraces mathematical formulae

² For a more detailed discussion of these points see Rausser and Freebairn [8].

based, for example, on Bayes' Rule for incorporating the additional information in an updated probability distribution function.

Summarizing, the policy problem is formulated as finding values for the policy variables to maximize the objective function subject to the current state of the system and probabilistic knowledge about its behaviour, and subject to the state transformation and probabilistic knowledge transformation functions. The derived adaptive policy strategies provide rational policy actions given the assumptions embodied in the optimization problem. Doubtful assumptions, e.g. concerning the objective function, may be parametrized to assess effects on, and the sensitivity of, policy strategies.

Policy Strategies

At least two characteristics of adaptive policy strategies are worthy of comment, namely, their feedback form and the dual nature of adaptive control.

The adaptive control strategies specify levels of the policy variables for each decision period as functions of information that would be available at the beginning of the period. That is, the policy action in each period is conditional on the state of the system at the beginning of the period and on the updated probability distribution for the uncertain terms of the problem. In this way adaptive policy strategies involve a process of sequential revision of policy actions in an ordered manner as new information becomes available.

One of the principal features of adaptive policy strategies is their dual nature. They embrace dimensions of control, of learning and of design of experiments, and in general the three need to be considered concurrently. The control dimension arises from direct effects of different levels of the policy variables in the objective function and from indirect effects on current and future levels of endogenous variables in the objective function. The learning dimension arises from the feedback form of the policy strategies and in particular from the process of using new information to update the probability distribution function for the uncertain terms of the policy problem. Directly and indirectly alternative policy actions influence the sample information generated by the economic system and this is the design of experiments dimension.

The dual nature of adaptive policy strategies has important implications for policy which would not be observed if the elements of uncertainty and sequential information accumulation were ignored. The inherent benefits of more information derive from a more precise perception of the effects of alternative policy actions which in turn raises the expected efficiency of future decisions. The information content of current policy actions may be enhanced by decisions which are not optimal in a short term control only context. For example, more can be learnt about the parameters of a supply function if observations are forced over a wide price-quantity surface than if a constant price is maintained, but the latter may result in a higher realization of the objective function in the short-term.³ In this situation

³ In a concrete context Dreze [5, p. 15] notes '... a monopolist may wish to depart from the price which maximizes expected profit simply to learn more about his demand function'.

the adaptive policy strategy may sacrifice current control efficiency in order to achieve greater future control efficiency via a more precise probability distribution function for the uncertain terms.

An Example

To illustrate the formulation of an Australian agricultural policy problem in an adaptive control framework, Australian wheat stabilization policy might be analysed in the following manner.⁴

The policy variables, variables that could be manipulated by the Australian and State governments, would include wheat delivery quotas, home consumption price, export guarantee price and level of guaranteed exports. Further, suppose these policy variables may be revised at the beginning of each (wheat) year. Alternative policy actions would consist of different combinations of levels of the policy variables.

Alternative policy actions would be evaluated in terms of performance variables considered to have welfare connotations for identifiable groups influential in the formulation of wheat stabilization policy. These groups include the Wheat Growers' Federation, the Treasury and Cabinet members. The preferences of these groups may be considered in terms of price received by the producer, stability of price over time, consumer price, Treasury outlay, and so forth. These variables become the argument variables of a set of objective functions. Different functions in the set would have parameters reflecting different relative weights between the argument variables. For example, a producer oriented function would place relatively more weight on higher producer price and less weight on lower Treasury outlay than would a Treasury oriented function.

Constraint functions delineating the policy possibility set would describe the endogenous variables of the objective function set as functions of the policy variables, other exogenous variables (e.g. seasonal conditions and the world wheat price) and lagged variables describing the current state of the wheat sector (e.g. carryover stocks and past prices). The constraint functions would be based on wheat supply and demand functions and on identities specifying producer receipts, Treasury outlays and other performance variables. In reality, policy makers would have imperfect knowledge about some of the parameters and variables of the constraint functions.

Probabilistic information about the unknown terms of the constraint functions could come from several sources. Econometric studies of the demand for and supply of Australian wheat provide information on the mean and variance-covariance terms of a probability distribution on the parameters of the functions. With respect to future values of the other exogenous variables, probabilistic estimates might be obtained, for example, for the world wheat price and for seasonal conditions from the Australian Wheat Board and the Bureau of Meteorology, respectively.

The above probability distribution functions would be updated as

⁴ Little imagination is required to characterize a diversity of contemporary Australian agricultural policy problems in the framework of an optimization problem. Practically speaking, the most important limiting consideration is likely to be the availability of an acceptable mathematical representation of behaviour of the economic system under study.

time proceeds. Additional sample observations would be used in re-estimating the econometric model of wheat demand and supply. The Australian Wheat Board would use forthcoming market outlook information to revise its estimates of future world wheat prices. These learning activities become the information updating functions of an adaptive control model.

The foregoing illustrates how Australian wheat stabilization policy might be formulated as a multiperiod optimization problem. Allowance is made for uncertainty about the effects of alternative policy actions. Forthcoming information may be used to learn about the uncertain terms as decision periods pass. The adaptive policy strategy is obtained by solving the optimization problem.

For each policy variable for each decision period a feedback function is obtained. Thus, for example, the policy variable wheat production quota is expressed as a function of variables describing the current state of the wheat economy, e.g. carryover stocks, and probabilistic information reflecting knowledge about behaviour in the system, e.g. of consumer demand, and probabilistic information about the other exogenous variables, e.g. world wheat price. At the beginning of each decision period the adaptive policy strategy together with updated information is used to compute the policy action for that period. For future periods the strategy indicates how policy actions are to be revised in the light of forthcoming information.

Linear-Quadratic Adaptive Control Problem

In an application context most of the literature on adaptive control theory, e.g. engineering-oriented studies by Aoki [1], Athans [2], and Bryson and Ho [3], and economic-oriented studies by Prescott [7], and Rausser and Freebairn [9], has focused on the linear-quadratic problem. This particular class of models requires assumptions which seem tolerable for a number of potential applications. This section provides a formal specification of the linear-quadratic problem and summarizes the available solution procedures and implied policy strategies.

Policy Problem

The linear-quadratic adaptive control model specification of a policy issue is represented as: find the adaptive decision strategy

$$(1) \quad u_t^*(y_{t-1}, P^{t-1}(\cdot)), t = 1, 2, \dots, T$$

to maximise

$$(2) \quad J = E \left\{ \sum_{t=1}^T D^{t-1} (2k_t' y_t + 2h_t' u_t - y_t' K_t y_t - u_t' H_t U_t - 2u_t' L_t y_t) \right. \\ \left. + D^T (2f_{T+1}' y_T - y_T' F_{T+1} y_T) \right\}$$

subject to

$$(3) \quad y_t = A y_{t-1} + B u_t + C x_t + e_t$$

$$(4) \quad P^t(A, B, C, x_t^T, e_t^T) = P_1^t(A, B, C) P_{21}^t(x_1) P_{22}^t(x_2) \dots P_{2T}^t(x_T) \\ P_{31}^t(e_1) P_{32}^t(e_2) \dots P_{3T}^t(e_T)$$

$$(5) \quad P^t(\cdot) = I_t(P^{t-1}(\cdot), y_t, u_t, x_t, s_t),$$

and

$$(6) \quad y_0 = y(0) \text{ and } P^0(\cdot) = P(0)$$

where y is an n -vector of endogenous (or state) variables, u is an m -vector of policy variables, x is a p -vector of other (noncontrollable) exogenous variables, and e is an n -vector of error terms. The particular objective function J in (2) is assumed to be concave, and its parameters in k, h, f, K, H, L, F and D are assumed known; E is the expectation operator. The parameter matrices A, B and C of the constraint functions (3) are assumed to be unknown.⁵ The function $P^t(\cdot)$ in (4) denotes the probability distribution function, or its set of sufficient statistics, conceived at the beginning of period t for the unknown terms, $A, B, C, x_t^T (= \{x_t, x_{t-1}, \dots, x_T\})$ and $e_t^T (= \{e_t, e_{t-1}, \dots, e_T\})$. The function $I_t(\cdot)$ of (5) denotes an information updating function where additional sample observations, y_t, u_t, x_t , and other new information bearing on x_t^T , denoted as s_t , is combined with $P^{t-1}(\cdot)$, to derive the updated probability distribution function $P^t(\cdot)$. The vector $y(0)$ denotes the initial state of the system and $P(0)$ denotes the initial probability distribution function. Before proceeding, a few comments on the underlying assumptions of the problem seem appropriate.

Even though the actual form of the objective function may not be quadratic, such a form might provide a reasonable approximation. Appeal could be made to a Taylor series expansion in which the linear and quadratic terms are retained. The function allows greater levels of an argument variable to increase (or decrease) social welfare at an increasing rate either dependent on or independent of the levels of other variables. For cases where a performance variable has a target path, a quadratic function can be specified to penalize deviations from this path with the penalties growing at an increasing rate.

Similarly, the assumption of linear equality constraints may be justified as a local approximation. Further, it is a common assumption in applied econometric studies. Since a difference equation of order r can be transformed into r difference equations of order one, the assumption of a set of first order difference equations is not restrictive.

For the purposes of this paper the probability distribution function $P^t(\cdot)$ is considered in terms only of its mean and variance-covariance statistics for the case where the constraints are based on an estimated econometric model. The latter supplies the respective statistics for the parameters A, B and C and for the error variables e_t^T . Note that an assumption of serial independence between the error terms is assumed. Forecasts of future values of the other exogenous variables x_t^T are assumed to be independent of the estimates of the econometric model.

With respect to the parameters in the matrices A, B and C , the learning function $I_t(\cdot)$ involves updating the econometric model estimates with the additional sample observations y_t, x_t and u_t . Potential updating procedures include Kalman filters, the application of a Bayesian regression estimator with an informative prior, and re-estimation by least squares using the augmented sample. In all cases the learning function will be a complex nonlinear function in $P^t(\cdot)$ and the additional sample observations.⁶

The adaptive control model specified in (2) through (6) admits a number of special models reported in the literature. The deterministic

⁵ This specification ignores uncertainty about the general form of the constraint functions, e.g. of appropriate included variables and of the mathematical form of the functions.

⁶ For further details see, for example, Aoki [1], Bryson and Ho [3] or Zellner [13].

problem follows if e_t is a null vector, and x_t , A , B and C are known. Holt [6], Theil [10] and others consider the case in which e_t and x_t are stochastic and A , B and C are known. Stochastic control problems discussed by Turnovsky [12], Chow [4] and others arise if $P^t(\cdot) = P(0)$ for all t .

Derivation of Adaptive Policy Strategies

At least conceptually the adaptive policy strategies $u_t^*(\cdot)$ in (1) may be obtained by dynamic programming or by the discrete maximum principle. Unfortunately, except for very simple problems it is impracticable to derive these strategies. The sources of intractability arise from the highly nonlinear form of the information learning functions (5), and from the "curse of dimensionality".⁷ Given the present state of knowledge, it is necessary to consider policy strategies which approximate the adaptive policy strategies.

Some Approximate Policy Strategies

Two approximate and hence suboptimal solutions to the optimization problem (1) through (6), namely the stochastic policy strategy and the sequential stochastic policy strategy, are discussed in this section.⁸ At the expense of further simplification of the policy problem these strategies are easily solved with the aid of a computer.

The stochastic policy strategy, denoted as $u_t^s(y_{t-1}, P(0))$, $t = 1, 2, \dots, T$, is obtained by finding the sequence of policy variables to maximise (2) subject to (3), (4) and (6) with $P^t(\cdot) = P(0)$ for all t . The latter means that the information updating function (5) of the adaptive control problem is bypassed. Now, the simplified optimization problem is the familiar linear-quadratic stochastic control problem. Operationally, $u_t^s(\cdot)$ is obtained by the backwards solution of a set of recursive equations. Details of the procedure may be found in Aoki [1] or Chow [4]. The stochastic policy strategies are in linear feedback form, *viz.*,

$$(7) \quad u_t^s(y_{t-1}, P(0)) = G_t(P(0))y_{t-1} + g_t(P(0))$$

where the $m \times n$ matrix G_t and the m -vector g_t are based on the known parameters of the objective function (1) and the mean and variance-covariance terms of the probability distribution function $P(0)$ in (6).

The stochastic policy action taken each decision period is computed using the new information on the current state of the system as measured by y_{t-1} . No provision is made for the use of forthcoming information to learn about the parameters A , B and C of the constraint functions. The sequential stochastic decision strategy incorporates learning about the A , B and C parameters.

The sequential stochastic policy strategy, denoted as $u_t^{ss}(y_{t-1}, P^{t-1}(\cdot))$, $t = 1, 2, \dots, T$, is obtained by artificially separating the control and learning aspects of the policy problem. Each decision period a new stochastic optimization problem is solved using the updated probability distribution function $P^t(\cdot)$; recall that for the stochastic policy strategy it was assumed that $P^t(\cdot) = P(0)$ for all t . To illustrate, after following the

⁷ For details see Aoki [1], Prescott [7] or Rausser and Freebairn [9].

⁸ Other approximate policy strategies are discussed in Aoki [1], Zellner [13] and Rausser and Freebairn [9]. They include linearization of the information learning functions and certainty equivalent policy strategy.

first period stochastic policy action, $u_1^s(y(0), P(0))$, the new information is used to compute $P^1(\cdot)$ according to the information learning function (5). Using $P^1(\cdot)$ a new stochastic optimization problem is solved to obtain $u_t^{ss}(y_{t-1}, P^1(\cdot)) = u_t^s(y_{t-1}, P^1(\cdot))$, $t = 2, 3, \dots, T$. The second period policy action, $u_2^{ss}(y_1, P^1(\cdot))$, is implemented, and the sequential process is repeated.

Compared with the adaptive policy strategy the sequential stochastic policy strategy recognises the control and learning dimensions but ignores the experimental dimension. In deriving $u_t^{ss}(\cdot)$ it is assumed that forthcoming sample observations will not be used to learn about the unknown A, B and C parameters, however this assumption is revised each decision period. That is, the sequential stochastic policy strategy only allows for the passive accumulation of information and ignores the influence of current policy actions on the information generated. It follows that the degree of suboptimality of the approximation will be related to the importance of the design of experiments dimension in the adaptive policy strategy.

The importance of the design of experiments dimension is reflected by two considerations, namely, the extent of imperfect knowledge and the extent to which alternative policy actions facilitate learning. The former is measured by the magnitude of the terms of the variance-covariance matrix $P^t(\cdot)$ (assuming unbiased estimates) and the latter is measured by reductions in these terms. The rate of learning is enhanced by policy actions which generate a wide dispersion of values for the explanatory variables of the constraint function. In a Monte Carlo study Prescott [7] found experimentation to be of little importance for cases where the ratio of the mean parameter estimate to its standard error exceeded one in absolute value. Further experiments would be desirable before regarding Prescott's results as generally applicable. Conservatively, it seems reasonable to conclude that for those applications based on econometric models with well-defined parameter estimates, that is, with parameter estimates significantly different from zero at the five per cent and higher levels, the sequential stochastic policy strategy will be a close approximation to the adaptive policy strategy.

Some Additional Results

The expected loss of alternative policy actions relative to the stochastic policy action and the expected value of additional information resulting in more efficient (i.e. lower variance) estimates of the unknown terms of the constraint functions can be analysed for the stochastic policy strategy, see, for example, Rauser and Freebairn [9]. With respect to the former, the expected first period loss of a policy action u relative to the stochastic policy action $u_1^s(y(0), P(0))$ is given by $(u_1^s(\cdot) - u_1)' N(u_1^s(\cdot) - u_1)$ with N a positive definite matrix. Characterization of the value of additional information, i.e. of more efficient estimates of the A, B and C parameters, of a model specification with smaller error term variances and more efficient forecasts of the other exogenous variables, provides a basis, together with the costs of these activities, for assessing the payoff of further refinements to the policy analysis.

An Example

This section provides an adaptive control model analysis of a sim-

plified version of one aspect of Australian wheat industry policy.⁹ Suppose the Australian government proposes to influence wheat production with the objective of attaining a desired stock of wheat and it has available one policy variable, the per tonne payment to producers, which may be revised each year, say, February of each year.

It is assumed that the policy objectives can be evaluated in terms of two variables, the stock of wheat available for the new marketing year, say at January 31, denoted as q_t (million tonnes), and the level of the producer price, denoted as u_t (dollars per tonne). For trade and storage reasons a desired wheat stock, q^* , is preferred. At current price levels increases in the producer wheat price raise social welfare but at a decreasing rate to reflect producers' pressure for increased returns modified by considerations of Treasury costs and domestic food prices. At least as a local approximation, the foregoing preferences can be represented in the framework of a quadratic objective function in the following manner

$$(8) \quad J = \sum_x^T (\alpha_1 u_t - \alpha_2 u_t^2 - \beta_1 (q_t - q^*)^2) \\ = \Sigma (2hu_t - u_t' H u_t + 2k'y_t - y_t' K y_t) + Q$$

where $y_t = [q_t : \cdot]$ is the state vector partitioned according to whether the state variables have a nonzero effect on the objective function or not,¹⁰ u_t and q_t are as defined above, and the parameters $h = 0.5\alpha_1$, $H = \alpha_2$, $k_1 = \beta_1 q^*$, $k_2 = 0$, $K_{11} = \beta_1$, and $K_{12} = K_{21} = K_{22} = 0$ remain to be determined.

Two objective functions and hence two sets of parameters for (8) are specified to reflect different weights attached to importance of the argument variables u_t and q_t in social preferences. Specifically, the parameters reflect the following assumptions: a desired wheat stock of $q^* = 13$ (million tonnes); a quasi-maximum wheat price of $u^* = 150$ (dollars per tonne)—a price beyond the likely observable level and so ensuring that for observable prices, higher prices are preferred—and two trade-off ratios between the argument variables u_t and q_t such that a million tonne divergence from the desired wheat stock causes about a half or about the same social loss as a twenty dollar drop in the producer price (from its quasi-maximum level).¹¹ The former trade-off reflects relatively stronger preferences for a higher producer price. The two sets of parameters are:

Objective Function A (strong preference for high price)

$$(8a) \quad h = 7.5, H = 0.05, k_1 = 6.5, K_{11} = 0.5, \text{ and:}$$

Objective Function B (weak preference for high price)

$$(8b) \quad h = 3.75, H = 0.025, k_1 = 6.5, K_{11} = 0.5.$$

⁹ Given the simplifying assumptions and the hypothetical data employed, the example is illustrative only.

¹⁰ Subsequently $y_t = [q_t : u_t]$.

¹¹ The estimated parameters are reached in the following way. Given that $J = 2az - bz^2 = a^2/b - b(z - z^*)^2$ with $z^* = a/b$ being the level of z at which J takes a maximum and the marginal condition $dJ/dz = -2b(z - z^*)$, then $b = -(dW/dz)/(2(z - z^*))$ and $a = bz^*$, where z denotes q or u . For function A, $dJ/du = dJ/dq = 1$, $q^* = 13$, $q = 12$, $u^* = 150$ and $u = 130$.

For the purposes of the policy analysis an inventory identity and a wheat supply function describe the important aspects of behaviour in the wheat sector and they form the basis of the state transformation functions. The inventory identity links the end of year wheat stock, q_t , to the beginning of year wheat stock, q_{t-1} , production, o_t , and sales, s_t ,

$$(9) \quad q_t = q_{t-1} + o_t - s_t$$

where all terms are in million tonnes. Wheat production, o_t , is specified as a function of the prices of wheat and wool, u_t and p_t , respectively, an index of seasonal conditions, r_t , lagged production, o_{t-1} , and an error term, ϵ_t , viz.

$$(10) \quad o_t = \phi_0 + \phi_1 u_t + \phi_2 p_t + \phi_3 r_t + \phi_4 o_{t-1} + \epsilon_t$$

where the ϕ 's are unknown parameters. After substituting (10) for o_t in (9), the state variables q_t and o_t may be specified in reduced form as functions of lagged endogenous or state variables, a policy variable, other exogenous variables and error terms, viz.

$$(11) \quad \begin{bmatrix} q_t \\ o_t \end{bmatrix} = \begin{bmatrix} 1.0 & \phi_4 \\ 0 & \phi_4 \end{bmatrix} \begin{bmatrix} q_{t-1} \\ o_{t-1} \end{bmatrix} + \begin{bmatrix} \phi_1 \\ \phi_1 \end{bmatrix} u_t + \begin{bmatrix} \phi_0 & \phi_2 & \phi_3 & -1.0 \\ \phi_0 & \phi_2 & \phi_3 & 0 \end{bmatrix} \begin{bmatrix} 1.0 \\ p_t \\ r_t \\ s_t \end{bmatrix} + \begin{bmatrix} \epsilon_t \\ \epsilon_t \end{bmatrix}$$

where all terms are as defined above. Equivalently, (11) may be restated as the following set of state transformation functions

$$(11a) \quad y_t = Ay_{t-1} + Bu_t + Cx_t + e_t$$

In the state transformation functions (11a), the parameters A, B and C and the variables in x_t and e_t are unknown. For an artificially generated sample of twelve observations ordinary least squares was applied to estimate the mean and variance-covariance terms of a probability distribution function for the ϕ 's and ϵ of (10)¹². The mean estimates together with the standard errors obtained are reported in (12) below.

$$(12) \quad o_t = 2.401 + .0527u_t - .0337p_t + .2389r_t + .6792o_{t-1} \\ (4.244) \quad (.0332) \quad (.0387) \quad (.2496) \quad (.3098) \\ E\epsilon_t^2 = .6702$$

The sample estimates in conjunction with the definitions of the elements of the parameter matrices of (11) are used to derive the mean and variance-covariance terms of the probability distribution function for the A, B and C parameters and for the e_t variables, e.g., $Ea_{11} = 1.0$, $Ea_{12} = 0.6792$, $\text{Var } a_{11} = 0$ and $\text{Var } a_{12} = (0.3098)^2$, where E and Var denote expected value and variance, respectively. The mean estimate

¹² Specifically, the data were generated as follows. The true form of the supply function was specified as $o_t = -1.5 + 0.045u_t - 0.02p_t + 0.25r_t + 0.85o_{t-1} + \epsilon_t$. The explanatory variables were generated from a normal distribution function with the parameters $u_t \sim N(100, 400)$, $p_t \sim N(150, 2500)$, $r_t \sim N(5, 4)$ and $\epsilon_t \sim N(0, 1)$.

of the vector of exogenous variables was arbitrarily specified as $E x_t = [1, 150, 5, 12]$ for all t . Collectively these statistics comprise the probability distribution function $P^0(\cdot)$,

$$(13) \quad P^0(A, B, C, e_t^T, x_t^T).$$

Summarizing so far, relations (8), (11) and (13) plus an initial state condition $y_0 = y(0)$ represent the wheat policy problem as a linear-quadratic stochastic optimization problem.

The derived stochastic policy strategy,

$$u_t^s(y_{t-1}, P^0(\cdot)) = G_t(P^0(\cdot))y_{t-1} + g_t(P^0(\cdot)),$$

for the two objective functions (8a) and (8b) is reported in Table 1. Because of the special assumptions of stationarity of $E x_t$ for all t the parameters of the policy strategy are time invariant in this example. The more interesting aspect of these strategies concerns the parameters of the G -vector. The greater the level of the present stock of wheat, q_{t-1} , and the greater the level of last year's wheat production, o_{t-1} , the lower would be the recommended producer wheat price for the coming season. Sensitivity of the recommended producer price to future levels of the state variables is less for objective function A which reflects a relatively stronger preference for high wheat prices. The respective first period stochastic policy actions for a current situation characterised by wheat stocks of fifteen million tonnes and last year's production of twelve million tonnes are shown in Table 1.

To illustrate the learning dimension of an adaptive policy strategy, assume five years additional information becomes available. The new sample observations can be incorporated in an augmented sample for re-estimating the unknown parameters of the wheat supply function (10). The updated estimated function obtained is

$$(14) \quad o_t = 1.918 + .0510u_t - .0330p_t + .2713r_t + .7074o_{t-1}$$

$$(4.107) \quad (.0138) \quad (.0194) \quad (.0998) \quad (.1263)$$

$$E \epsilon_t^2 = .7934$$

Comparing the estimates of (12) and (14), inclusion of the additional sample observations in the estimates has improved the efficiency of the parameter estimates markedly. The updated estimates, together with additional information about the other exogenous variables, e.g. suppose

TABLE 1
Stochastic Policy Strategy and First Period Policy Action for Two Objective Functions

Objective Function ^a	Parameters ^b			Policy Action for $y(0) = [15, 12]$ (producer wheat price)
	G-vector with respect to:		g-scalar	
	q_{t-1}	o_{t-1}		
Function A	-2.119	-3.419	178.8	105
Function B	-2.871	-4.302	186.2	92

^a Parameters given in (8a) and (8b) respectively.

^b Conditional on $P^0(\cdot)$ in (13).

$Ex_t = [1, 120, 5, 10]$, is collated to form the updated probability distribution function $P^4(\cdot)$. A new stochastic optimization problem using $P^4(\cdot)$ rather than $P^0(\cdot)$ is solved to derive the sequential stochastic policy strategy

$$u_t^{ss}(y_{t-1}, P^4(\cdot)) = G(P^4(\cdot))y_{t-1} + g P^4(\cdot).$$

The parameters of this strategy are reported in Table 2.

TABLE 2
Sequential Stochastic Policy Strategy for Two Objective Functions

Objective Function ^a	Parameters ^b		
	G-vector with respect to:		g-scalar
	q_{t-1}	o_{t-1}	
Function A	-2.302	-3.668	150.5
Function B	-3.190	-4.685	160.9

^a Parameters given in (8a) and (8b), respectively.

^b Conditional on $P^4(\cdot)$ described in text.

Conclusions

The paper explains how adaptive control methods can be applied to the analysis of policy making in the agricultural sector. Policy making is characterized as a multiperiod optimization problem in an uncertain world. A process of sequential decision making permits the use of forthcoming sample and other information to learn about the uncertain elements as decision periods pass. Adaptive policy strategies recognize the effects of alternative policy actions in terms of control, learning and design of experiments dimensions. To a large degree an adaptive control model assists in formalizing many of the assumptions which are implicit in a real world policy making situation.

Much of the literature has focused on the linear-quadratic model, partly because of the mathematical tractability of this class of optimization problems and partly because the required assumptions appear reasonable approximations for many potential applications. Inexpensive computer routines are available for deriving the stochastic and sequential stochastic policy strategies.

The adaptive policy strategy indicates policy actions for the current period and how policy actions for future periods should be revised in the light of changing circumstances and new knowledge about behaviour of the system. Parametric analyses using different objective functions, different economic model specifications and different forecasts of other exogenous variables could provide useful insights into the effects of alternative policy actions.

References

- [1] Aoki, M., *Optimisation of Stochastic Systems*, New York, Academic, 1967.
- [2] Athans, M., 'The Discrete Time Linear-Quadratic-Gaussian Stochastic Control Problem', *Annals of Econ. and Soc. Measurement*, 1: 449-491, 1972.
- [3] Bryson, A. E. and Y. Ho, *Applied Optimal Control*, Waltham, Mass, Blaisdell, 1969.
- [4] Chow, G. C., 'Effect of Uncertainty on Optimal Control Policies', *Int. Econ. Rev.*, 14 (3): 632-645, 1973.

- [5] Dreze, J. H., 'Econometrics and Decision Theory', *Econometrica*, 40(1): 1-18, 1972.
- [6] Holt, C. E., 'Linear Decision Rules for Economic Stabilisation and Growth', *Q. J. of Econ.* 17(1): 33-42, 1973.
- [7] Prescott, E. C., 'The Multi-Period Control Problem Under Uncertainty', *Econometrica*, 40 (6): 1043-1058, 1972.
- [8] Rausser, G. C. and J. W. Freebairn, 'Estimation of Policy Preference Functions: An Application to U.S. Beef Import Policy', *Rev. of Econ. and Stats.*, November, 1974.
- [9] ———, 'Approximate Adaptive Control Solutions to U.S. Beef Trade Policy', *Annals of Econ. and Soc. Measurement*, 3: 177-204, 1974.
- [10] Theil, H., *Optimal Decision Rules for Government and Industry*, Amsterdam, North Holland, 1962.
- [11] Tinbergen, J., *On the Theory of Economic Policy*, Amsterdam, North Holland, 1964.
- [12] Turnovsky, S. J., 'Optimal Stabilisation Policies for Deterministic and Stochastic Linear Economic Systems', *Rev. of Econ. Studies*, 60(1): 79-95, 1973.
- [13] Zellner, A., *An Introduction to Bayesian Inference in Econometrics*, New York, Wiley, 1971.