Agricultural response analysis in a longer term framework

R.J. Farquharson [*] , **O.J. Cacho** [**] **and J.E. Turpin** [***]

Abstract

Issues of long term soil fertility decline and sustainability are becoming more
important for cropping industries in Australia. Helping to manage the level of soil
fertility in this context is an aim of economic response analysis. This paper reviews
the theory and methods used by economists to derive the optimal level of an input to
be used in a production process. In particular, response functions generated by a crop
simulation model are used as a basis for the analysis. The use of such models is
becoming widespread in the research and extension community. A variety of
methods are presented, in increasing order of complexity, to account for the real
world characteristics of the production environment in this context.

*   NSW Agriculture, Centre for Crop Improvement, Tamworth and contact
        for correspondence: bob.farquharson@agric.nsw.gov.au
**   University of New England, Armidale
***   NSW Agriculture and APSRU, Toowoomba

1. Introduction

An important issue for grain growers in Australia is the management of soil fertility through a spectrum of management decisions, including artificial fertiliser applications and crop-sequence choices, in situations of long-term soil fertility decline. In areas of historically fertile soils, where past cropping practices have often not included the addition of fertilisers, the question is becoming one of sustainably managing soil fertility levels given a current level of infertility.

Two broad methods of changing soil fertility are possible - (1) the addition of artificial fertilisers, and (2) changing crop-fallow sequences (including the addition of legumes and altering fallow length). These decisions are made not only to manage soil fertility levels, but they also impact issues such as soil erosion, or control of crop pests, weeds and diseases. It is likely that a combination of these methods will be used in many cases. These decisions are being made in a farming systems context, because of the potential interdependencies between components of the system.

Whichever crop management is used, the practice of adding some artificial fertiliser to an existing level of soil fertility will be an important option when growing a crop. That there is a positive response to fertiliser is generally accepted, the question then becomes one of how much fertiliser to add.

The aim of this paper is to review the methods and underlying theory available to analysts in determining answers to the 'how much' question, accounting for the important economic, biological and climatic factors bearing on the decision. These methods are illustrated through the analysis of a synthetic crop-response dataset. An important question that arises is whether the incorporation of more economically-oriented prescriptions leads to a significantly better outcome than other methods which use less-sophisticated calculations.

The paper is organised as follows. Section 2 contains a review of the classical economic method of determining 'best operating conditions' and emphasises the assumptions underlying this approach. Section 3 outlines a number of cases where these classical assumptions break down and presents methods to overcome each problem. Section 4 details a set of data that has been derived to illustrate the methods discussed. Then in Section 5 the response analyses are presented, in increasing order of complexity, to account for the exceptions to the classical assumptions. Finally a conclusion and implications for analysts are presented.

2. The economics of response analysis

Anderson, Dillon and Hardaker (1977) set out the principles of decision-making based on rational choice. In making decisions, especially those characterised by the presence of uncertain outcomes based on random factors (ie stochastic outcomes), the decision-analysis paradigm is proposed as an aid to decision makers. When the consequences of decisions are important, and when there is uncertainty about which consequence will occur, a rational consideration of the available information will

allow a systematic approach to making the best possible decision, according to the objective of the decision maker.

The decision-making process can be characterised as a 'choice-chance-consequence' process. Faced with a set of information and uncertainties, the decision-maker has available a number of alternative acts or choices, which involve potentially a number of different events (outcomes or states of nature), each with a particular chance of occurring. There may also be quantifiable consequences for different management choices, depending on which state occurs. These outcomes may be expressed as payoffs (in dollar terms). The decision-maker has certain objectives, or criteria, by which decisions will be made. Given all of this, the process of decision analysis involves systematically quantifying these alternatives as an aid to making the best possible decision.

Gaining further information (such as purchasing a forecast, or undertaking more research) can be worthwhile if it refines the probabilities of outcomes so that a different decision is made. This mode of analysis can be used to develop a strategy, or "a recipe for future action conditional on observed experimental outcomes" (Anderson, Dillon and Hardaker p. 6). Thus the management decision can be made according to some observed interim outcome.

Two general farm management cases where economic principles apply are the selection of an input level within an enterprise and the selection of enterprise mix to use in a whole-farm context. Makeham and Malcolm (1981) emphasise the importance of whole-farm analysis in determining management recommendations about choice of enterprises (crops, pasture and livestock mixes) when there are a number of alternative enterprises. In this paper the focus is on selection of a level of input to use, which is considered here separately from the question of enterprise selection.

2.1 Deterministic single output

The classical approach to agricultural decision analysis has been outlined by, among others, Anderson (1967) and Dillon and Anderson (1990). Dillon and Anderson discuss a normative approach to solving the problem of manipulating an input to a production process to determine the best operating conditions to achieve a specific goal. If profit maximisation is considered an important aim for commercial grain growers, then the likely biological response of the crop to the input, given the relevant input and output prices, can be used to determine the 'best', or most profitable, level of input to use.

In the case of certain (ie deterministic) outcomes, the biological response function for a single output can be represented as:

$$Y = f(X_1, X_2, ... X_n),\tag{1}$$

where Y is the quantity of output and the X's are the quantities of inputs.

France and Thornley (1984) have characterised this function as being generally of the diminishing returns type. Similarly, Dillon and Anderson base their theory of crop and livestock response on three assumptions about relationship (1). These are:

1. that there is a continuous smooth causal relationship between the inputs and the output, implying that the first derivative of the response function (1), $\partial Y / \partial X_i$, exists;
2. that diminishing returns prevail with respect to each input factor X_i, implying that for each input $\partial Y / \partial X_i$ is positive but decreasing as X_i increases, and that the function is concave so that the Hessian matrix of second partial derivatives is negative definite (Simon and Blume 1994); and
3. that decreasing returns to scale prevail, so that an equal proportionate increase in all inputs results in a less than proportionate increase in output, implying $\sum (X_i / Y)(\partial Y / \partial X_i) < 1$.

Dillon and Anderson state that "there is no strong evidence of any but diminishing returns" under the usual conditions of crop and livestock production.

Assuming no constraints on inputs or output, they then write the profit function as:

$$\pi = p_y Y - \sum_i p_i X_i \qquad (p_y, p_i \geq 0) \quad (i = 1, \ldots ,n) \qquad (2)$$

where π is the net gain from the response process, or profit, and p_y and p_i are the (fixed) prices of the output and inputs, respectively.

By deriving the first order conditions and assuming sufficiency holds, Dillon and Anderson state the necessary conditions for the profit maximising input levels to be:

$$p_y * \partial Y / \partial X_i = p_i. \qquad (3)$$

Solving these n equations yields the set of $X_i's$ that constitute best operating conditions, or the optimum levels of inputs to use. Condition (3) can be interpreted as requiring that each input be added until the marginal value product equals the price of the input. Alternatively the condition is that the slope of the production function (marginal product) equals the price ratio of input to output. The input increments are added while ever extra revenue exceeds extra costs, and stops when marginal revenue equals marginal cost. This is the marginalist paradigm that is the basis of the neo-classical approach to economic analysis.

For the one-input case, the solution (3) is illustrated in Figure 1. For the response or production function, $Y = f(X)$, the economically optimum (or best) level of X to use, X_e, is found where the slope of the production function, the marginal productivity $\partial Y / \partial X$, equals the price ratio p_i / p_y. Note that this is different from the point of maximum yield X_y, which should not be used to determine the economic input level because it does not account for the incremental cost of inputs relative to the value of output.

One way of determining the best level of input to use (ie best operating conditions) is to calculate enterprise gross profit, expressed as the Gross Margin (GM), for different levels of the input. The GM is defined for a production unit (eg one hectare of land sown to wheat) as the unit returns less variable costs. When GM, which is equivalent to π above, is plotted against an input, the point of maximum profit is taken to be the best level of the input to use. By definition the maximum point of this relationship corresponds with the optimum level of input to use to maximise profit, as derived above using (3).

This profit function is relatively easy to calculate and plot when only one input level is being considered, but is difficult for two and impossible for three or more inputs considered simultaneously. In this latter case, the algebraic approach in (1) to (3) above must be used to determine best operating conditions.

One aspect of many profit functions calculated in this way and plotted from experimental data is that they are relatively flat or unresponsive for levels of input near to the profit maximising point (Anderson 1975).

2.2 Deterministic multiple output

Dillon and Anderson also consider the case of unconstrained multiple responses, or the simultaneous production of a variety of outputs. Examples that they give are grain and straw in cereals, muscle and fat in livestock or various grades of meat in beef production. In these cases it is impossible to allocate inputs between the various responses, that being done within the plant or animal. They posit r simultaneous responses, each characterised by a response function:

$$Y_k = f_k(X_1, X_2, ..., X_n) \qquad (k = 1, ... r), \qquad (4)$$

and write the profit function as:

$$\pi = \sum_k p_k . Y_k - \sum_i p_i . X_i \qquad (i = 1, ... n). \qquad (5)$$

Subject to the sufficiency condition that the Hessian matrix be negative definite, best operating conditions can be derived from the simultaneous solution of the set of n equations:

$$\sum_k p_k . (\partial Y_k / \partial X_i) - p_i = 0. \qquad (6)$$

2.3 Stochastic single output

Dillon and Anderson (1990) and Anderson, Dillon and Hardaker (1977) also consider the derivation of best operating conditions under conditions of variable (ie. stochastic) or risky outcomes. This involves maximising a utility objective function for the decision maker which is directly proportional to profits, these being in turn related to the response function Y and prices, which may reflect uncertainty. If decision maker is risk averse, as is generally assumed, the best operating conditions derived under

risk provide for a lower level of input than the deterministic case above, because variance of income constitutes a friction to efficient production or a risk-induced increase in input costs.

Anderson, Dillon and Hardaker (Chapter 6) derive a modified version of (3), which states that the optimal level of input occurs when the marginal factor cost (or input price) is equal to the value of the marginal expected product minus a marginal risk deduction that depends on the utility function and the marginal variance of revenue. Alternatively the best input level is where marginal revenue is equal to marginal cost (with respect to changing expected output) plus a marginal addition due to risk.

3. Exceptions to these cases

For wheat growing in the North Eastern cropping region of Australia, the deterministic or riskless theory above breaks down in a number of important ways. One of those ways involves issues of variability, but there are others that arise for other reasons. These cases and possible remedies are discussed in the next sub-sections.

3.1 Yield and grain quality response and non-fixed prices

The first of these cases involves the assumptions of smoothness and concavity of response and price. The wheat response to added artificial fertiliser nitrogen can be expressed in terms of two outputs, yield (quantity) and protein content (quality). The generalised response relationship presented by Strong (1981) is shown in Figure 2. The yield response is of the form of (1) and accords with the standard assumptions stated above. However, over at least part of the domain of added nitrogen, the protein content is hypothesised to be a convex response to the added input. The first order conditions can still be derived if the response is smooth and differentiable, however the second order sufficient conditions need to be checked for a maximum.

In addition, prices received for wheat are an increasing and non-smooth function of protein content. Figure 3 shows prices at two different times in 1997. The price per tonne is set according to grade (based primarily on protein content) and protein increments within grade. Wheat grades (such as Australian Standard White, Australian Hard and Prime Hard) are based on protein levels plus other criteria such as screenings, grain hardness, milling quality and dough properties. Protein content has a greater influence on overall processing quality than any other single factor. The stepped nature of the price function has implications for incentives to wheat growers and nitrogen application decisions (Fraser 1998).

The stepped and non-smooth price function causes the classical economic response analysis to break down, because of its reliance on smooth functions and algebraic derivation of first-order conditions.

Economists have also investigated the use of Linear Response and Plateau (LRP) functions to represent the yield response to added inputs (eg Lanzer and Paris 1981, Waggoner and Norvell 1979). This notion is based on von Liebig's 'Law of the Minimum', which states that plant growth is proportional to an increase in the supply

of the most limiting factor (the "minimum factor"), see Figure 4. Jomini et al. (1991) show an example of a response analysis using such a specification. LRP functions also violate the classical smooth function assumption.

3.2 The multiple non-separable output case

The second case involves the characteristic of wheat response to added nitrogen whereby there are two outputs (yield and protein content). For wheat, the response function cannot be written as in (4), because the outputs are jointly determined by the inputs. So the function must be written as:

$$(Y_1, Y_2) = g(X_1, ..., X_n). \tag{7}$$

In the case of wheat, outcomes of yield and protein content are not separable in terms of price received. Payment is made on a unit basis that is determined by yield (ie $/tonne), but based primarily (although not exclusively) on quality characteristics.

Anderson, Dillon and Hardaker (1977) define an objective function that is separable as one that can be expressed as the sum of separate functions of single variables. If (7) was to be used in a profit function such as (5), it would be separable if the yield and protein effects were independent and could be simply added together. Importantly, Dillon and Anderson's profit function (5) requires that both outputs have a separate price.

Yield and protein are jointly determined by seasonal characteristics. Martin et al. (1996) illustrate how conditions during the grain filling period affect grain protein percentage. For a similar amount of total nitrogen in the grain harvested, yield (in terms of starch) and quality (protein) vary inversely according to whether the finish to the season is 'hard' or 'soft', that is, according to the temperature and moisture conditions during the grain filling stage. If temperature is relatively high then, when moisture becomes limiting due to high evapo-transpiration, the plant partitions energy more into protein at the expense of starch. When temperature is lower, or moisture is less limiting, the plant places more emphasis on starch production (higher yield) and less on protein.

Therefore, (7) should be written as:

$$(Y_1, Y_2(Y_1)) = g(X_1, ..., X_n), \tag{8}$$

and Dillon and Anderson's best operating conditions for the multiple response case do not apply because of the non-separable wheat response to added nitrogen.

3.3 Accounting for non-smoothness and multiple non-separable outputs

One way to deal with the problems of non-smooth responses and prices, and the non-separable joint response, is outlined in CIMMYT (1988). In making good recommendations regarding evaluation of technologies from the farmer's point of view, three premises are made about decision-making:

1. farmers are concerned with the benefits and costs of particular technologies;
2. they usually adopt innovations in a stepwise manner; and
3. they will consider the risks involved in adopting new practices.

Leaving aside the third consideration for the time being, the first two can be incorporated by the use of partial budgets, a marginal rate of return analysis and the target rate of return concept.

Partial budgeting involves determining costs that vary between input levels, principally the application of extra nitrogen and harvesting costs for the resulting higher-yielding crops. Gross benefits result from the extra yield and/or protein levels, and net benefits are the difference between extra benefits and extra costs. A net benefit curve can be plotted and the marginal rate of return can be calculated as the change in net benefits divided by the marginal costs, expressed as a percentage. Treatments that are dominated by others (always inferior) can be deleted from the analysis.

A target or minimum rate of return needs to be determined as a criterion for whether funds should be invested incrementally in the input or technology. Investment continues while the marginal rate of return remains above the minimum or target rate. CIMMYT (1988) discusses how the target rate of return might be determined. It could be set relatively high, eg 50% or 100% per annum for an adjustment to existing farming practices or implementing a new technology respectively. Alternatively, a rate of twice the cost of capital might be used. This rate might be in the order of 20 to 25%. Extension officers sometimes refer to the '2 to 1' rule which, if referring to a $2 return for $1 outlay, corresponds to a 100% rate of return.

Note that this approach does not depend on concavity, continuity or smoothness assumptions, and also overcomes the problem of joint and non-separable responses. But it does retain the marginalist approach and additionally includes the concept of opportunity cost. The equalisation of marginal benefits and marginal costs in (3) now includes as a cost the required rate of return on the invested capital. Therefore this is a more general approach than calculating the economic response (ie, GM) over input levels because it incorporates a specific required rate of return.

It is important to understand the rationale for setting a target rate of return. In CIMMYT (1988) working capital is defined as the value of input allocated to an enterprise with the expectation of a return within the production period. The cost of this capital is the benefit given up by tying the capital into the enterprise for that period of time. This may be a direct cost (if borrowed) or an opportunity cost (if earnings are given up from the best alternative use). For a farmer to invest in an input, the need is to consider the level of additional returns beyond the cost of capital that will satisfy some criterion that the investment is worthwhile. Something must be added to the cost of capital to repay the farmer for the time and effort in learning to manage, or trying, a new technology. CIMMYT (1988) does not interpret the extra return as a return to offset risk, although the uncertainty in agronomic and economic data is mentioned as an issue. Risk is addressed in the next section.

The target rate of return concept has been incorporated into the determination of classical best operating conditions be Jauregui and Sain (1992) for a single output response function. Adapting their notation to that used here, the best operating conditions are given by:

$$p_y . \partial Y / \partial X_i = p_i . (1 + R) \tag{9}$$

where R is the minimum rate of return acceptable to the farmer. Note that in Figure 1, the solution to (9) will lead to a lower level of X_e than the level given by (3), because the slope of the price ratio line is greater. The extra return required to cover the cost of capital and other factors included in the target rate of return leads to a higher marginal revenue requirement to cover marginal costs.

The target rate of return concept could be included in the plot of the profit (or GM) function against fertiliser input, as shown by Jauregui and Sain. Adapting their notation, the profit or GM would need to be adjusted to account for the target rate of return at each level of input. The profit function of (2) becomes:

$$\pi = p_y . Y - p_n . N . (1 + R) - TCNV \tag{10}$$

where N incorporates the total costs that vary and $TCNV$ denotes the variable costs that do not vary as N is increased.

3.4 Variability in response and price

The classical approach of Dillon and Anderson includes the case of variable outcomes with concave response functions. A complicating case to consider is where variability in yield, protein and price is allied with non-smooth and non-concave functions and the multiple non-separable issues mentioned above. In the northern cropping region of NSW and Queensland the crop growing environment is characterised by substantial climatic variability with respect both to in-crop rainfall and the moisture and temperature conditions during grain fill after crop tillering, and to fallow conditions. An example of the response of yield and protein to different levels of nitrogen input in different seasons is presented later in this paper.

In addition there is price risk. Figure 3 shows that there can be a wide divergence in prices, and prices can change on a weekly basis. At sowing the extent of subsequent price movements and payment schedules may be unknown. However, hedging and forward contracting options are now available to wheat growers which reduce the importance of price risk. Hence, the issue of variability in this analysis has only been considered with respect to the impact of climate on yield and protein outcomes.

An alternative, and more general, way of considering the best operating conditions via the distribution of returns is through comparing entire probability distributions according to specific rules. The essence of these approaches is in determining whether particular cumulative distribution functions (CDFs) are dominated, or whether certain distributions (even when they cross) might be preferred by a risk-averse decision-maker.

3.4.1 Methods of comparing distributions

Anderson, Dillon and Hardaker (1977) have espoused the use of expected utility theory in representing the choices of decision makers. That is, Bernoulli's principle, which is deduced from a number of axioms, states that a utility function exists for a decision maker whose preferences are consistent with those axioms. This function associates a single real number (utility value) with any risky prospect. When utility is plotted against income, the utility curve is posited to be concave if the higher levels of income are associated with increased riskiness or variation in income. The degree of risk aversion is associated with the curvature of this function.

There are three common measures of risk aversion, the absolute, partial and relative risk aversion coefficients. In Bardsley and Harris' (1987) terminology, these are respectively:

$$\alpha \equiv -U''(W) / U'(W)$$

$$\beta \equiv -[U''(W) / U'(W)].W$$

$$r \equiv -[(U''(W_0 + \pi) / U'(W_0 + \pi)].\pi .$$

Here W is average wealth of the decision maker, W_0 is initial wealth, and π is stochastic income. It can be noted that although these measures are generally specified in terms of wealth levels (a stock concept), they have often been applied to distributions of income or profit (which are flow variables).

The measure of absolute risk aversion, α, is not unit-free, ie. it depends on the level of wealth. The measures of relative and partial risk are both unit-free. When using a particular measure of risk aversion to compare income distributions, and in considering parameters from other studies to make comparisons, it is desirable to utilise a unit-free measure (Bardsley and Harris).

In comparing distributions of outcomes, first degree stochastic dominance (FSD) and second degree stochastic dominance (SSD) rules can be applied. These decision rules imply that marginal utility is positive $(U'(W) > 0)$, and that it is positive and diminishing $((U'(W) > 0)$ and $(U''(W) < 0))$, respectively. The main problem for SSD is the case of low crossing points of the CDFs, since SSD always requires the dominant distribution to have a greater minimum than the dominated distribution because stochastic dominance rules assume that the individuals belong to the class of all risk averters including infinitely risk averse individuals. The question for cases of low crossings of CDFs can only be resolved based on the particular level of risk aversion of individuals.

Parton and Carberry (1995) discussed the theoretical and practical shortcomings of stochastic dominance and mean-standard deviation analysis. These included the assumption of risks being nonnormally distributed. Just and Weninger (1999) have recently cast some doubt on a number of analyses showing yield distributions to be nonnormal.

The extension developed to overcome this problem of low crossing points of CDFs involves placing bounds on the risk aversion parameter and determining the outcomes when the parameter fall into particular numerical ranges. Two techniques have been developed to overcome the low crossing problem. The first is generalised stochastic dominance (GSD), which starts from the expected utility function. Meyer (1977) investigated the magnitude of this expression under the conditions that the Pratt (1964) risk aversion coefficient (α in the terminology above) falls within an interval

$$\alpha_1 \leq [\frac{U''(W)}{U'(W)}] \leq \alpha_2 .$$

Using this framework, the approach is to look at the utility difference between two distributions but hold the risk aversion parameter in a particular interval. Meyer solved this problem using an optimal control framework, so that when the decision maker's utility function has $\alpha(W)$ in the interval $[\alpha_1, \alpha_2]$ then a particular distributional dominance holds. This is a numerical evaluation technique for given probability distributions, a computer program (MEYEROOT) is available for this analysis. The problem with this approach is the use of arbitrary α_1 and α_2 values. It should be noted that GSD is a generalisation of other stochastic dominance forms. Parton and Carberry demonstrated the use of generalised stochastic dominance (or stochastic dominance with respect to a function) for a simulation model dataset of kenaf production in Northern Australia.

An alternative approach was described by Hammond (1974) who showed that, for two alternatives whose CDFs cross once, under constant absolute risk aversion there is a break-even risk aversion coefficient (BRAC) that differentiates between these two alternatives. Anyone with a risk aversion coefficient (RAC) larger than the BRAC will prefer one alternative, while anyone with an RAC smaller than the BRAC would prefer the other. The important difference between this technique and Meyer's GSD is that rather than specifying RAC bounds, one can solve for the BRAC and then proceed to investigate whether it is reasonable for individuals to have RACs that are larger or smaller than the particular value. However, this requires knowledge of the functional form of the assumed probability distribution.

McCarl (1988, 1990) developed a computer program (RISKROOT) to implement Hammond's approach with an empirical discrete distribution of unknown form. The program takes data for two alternative distributions and searches for the BRAC(s) which distinguish between them.

Of the techniques that can be used to resolve stochastic dominance choices (ie. Meyer's GSD using MEYEROOT and McCarl's RISKROOT), the latter has the advantage that it identifies the BRAC points at which preferences switch. This gives stronger results than GSD.

3.4.2 Interpreting results in terms of Risk Aversion Coefficients

The RISKROOT program generates results showing which distributions are dominant for ranges and values of RACs. As discussed above, these RACs are the Pratt absolute risk aversion coefficient. When interpreting the results of such comparisons for groups of decision makers (eg farmers), an important question relates to the levels of risk aversion that those decision makers might usually exhibit, and therefore which choice they would be most likely to make. From the viewpoint of the R&D institution or researcher trying to interpret experimental or other results, the important question is: "Which outcome is most likely to be chosen by a risk-averse profit-maximising farmer or group of farmers"?

This question is necessary when making recommendations to farmers based on R&D activities for decision support within extension or advisory programs. Vital information in this context is evidence of the typical levels of risk aversion for farmers within the target group of decision makers.

What is an appropriate level of risk aversion for Australian farmers? There have been a number of studies of this question, the most recent and comprehensive is by Bardsley and Harris. Bar-Shira, Just and Zilberman (1997) estimated risk aversion parameters for Israeli farmers and compared them to the Bardsley and Harris results for Australian farmers.

Bardsley and Harris estimated α at the median level of mean income and wealth for the wheat-sheep zone to be $2.1x10^{-5}$. Patten, Hardaker and Pannell (1988) derived a range of values for α of $6x10^{-6}$ to $2x10^{-6}$. Parton and Carberry used a figure of $3x10^{-5}$. Bar-Shira et al. estimated the median value of this coefficient for Israeli farmers to be $4.4x10^{-6}$. Patten et al. considered the appropriate range for this parameter to be between $1/W$ and $3/W$ (where W is net wealth as defined above).

If it is considered desirable to use a unit-free measure of risk aversion, then the equivalent relative risk aversion coefficient parameter (r) should be derived. Patten et al. suggested that a plausible range of r from the literature is between 1 and 3. Bardsley (personal communication 1999) has suggested the range is from 0.5 to 2. Bar-Shira et al. estimated the median value of this parameter to be 0.615 for Israeli farmers. Average farm net worth over three years (1995-97) in the northern NSW Wheat Sheep Zone is $1.27 million (ABARE 1999). Using a range of the relative risk aversion parameter from 0.5 to 3 together with this figure, α is estimated to lie in the range $3.9x10^{-7}$ to $2.4x10^{-6}$.

Therefore, working from a generally-accepted relative risk aversion coefficient range and a relevant measure of current farm wealth for the target population of farms, the appropriate level of α to use as a BRAC comparison in the RISKROOT output has been derived as $3.9x10^{-7}$ to $2.4x10^{-6}$. This is quite comparable with values estimated directly by other authors.

3.5 Implications of the Nitrogen cycle

The last, and perhaps the most important, divergence from the classical static response analysis is the case of accounting for the effects of the changing stocks of fertility in the soil over time, due to the effects of the nitrogen cycle. Hayman and De Vries (1995) have illustrated the nitrogen cycling processes in a natural (balanced cycle) and a continuous cereal (broken cycle) systems. Organic soil nitrogen is a relatively large pool from which the mineral (plant available) nitrogen is derived by mineralisation. In cropping systems with a net export of nitrogen, only small amounts of mineral nitrogen are returned (immobilised) to the organic pool, resulting in decline of the soil nitrogen fertility.

Crop land fertility decline over time is the process of long term decline in nitrogen in the soil. Growing a crop without adding fertiliser means that after grain protein is removed at harvest there will be less soil and plant available nitrogen for the next crop. Continuation of this process means that crop yield and/or protein levels will fall over time. Therefore, both the organic and mineral nitrogen levels in a soil will decline with repeated cropping in the absence of added fertiliser nitrogen. This has been the experience under long cropping phases in many areas of initially-fertile soil in Australia (Dalal and Mayer 1986).

A static (or single-period) response analysis does not account for this carryover characteristic, especially when it is combined with a stochastic environment. The incorporation of the dynamics of nitrogen accumulation and loss with exploitative land use and climatic effects over time addresses this issue.

The amount of organic nitrogen is associated with the level of organic carbon in the soil. Average levels of fertility decline through continuous cropping of a black vertisol soil over a long period can be hypothesised. Levels of organic carbon in Waco soils of southern Queensland have declined by 28%, and levels of mineralizable nitrogen have declined by 55% as a result of this fertility decline in Queensland (Dalal and Mayer 1986).

In this paper we are interested in the nitrogen fertility directly available to the wheat plant, ie in mineral rather than organic nitrogen (even thought the two are linked). We will use the term 'soil fertility' to refer to the mineral nitrogen which is made available to the plant via application of artificial nitrogen. The soil fertility levels that have been generated in this paper are fertiliser nitrogen plus soil-derived mineral nitrogen.

3.6 Accounting for the nitrogen cycle

The existence of nitrogen effects in the next year due to actions in the current year means that the static or single year decision must be extended over a longer period.

3.6.1 Nitrogen budgeting

One way to account for the nitrogen cycle in deriving farmer recommendations for fertiliser application to wheat crops is to utilise a nutrient budgeting approach (Martin et al. 1996). When grain is harvested, nitrogen (protein) is also harvested and removed from the system. The total amount removed can be estimated from the

actual or projected yield and protein content. Then the required amount of nitrogen is estimated according to what will be removed, the mineral nitrogen demand needs to account for efficiency of uptake (although this depends on the protein content of the final crop to some extent). The soil nitrate supply can be determined by use of soil tests, or in some cases by interpreting historical yields and protein contents. Account must also be made for late fallow and in-crop mineralisation or release of nitrate in the soil from the pool of mineral nitrogen. Then the difference between demand and supply can be determined, and the amount of fertiliser to apply can be calculated.

This is a useful process to achieve a steady-state in soil fertility, but the approach does not account for economic factors. And it uses a target yield approach rather than an optimising method. It must be remembered, though, that the nutrient budgeting approach is applied as a learning tool for farmers and that it does stress the importance of the concept of nutrient balances in making cropping decisions. Also it is often difficult to be precise when the underlying data sources are not exact and the climatic effects are important.

Farmer practices in fertiliser application for wheat appear to involve using either zero fertiliser, a fixed annual rate (a rule of thumb) or a tactical (flexible) approach on a year to year basis (Hayman and Alston 1999). The extent to which farmers believe and utilise results from R&D simulation models of fertiliser response is uncertain (Hayman and Alston). However, with the aim of evaluating alternative tactical approaches to nitrogen fertilisation, the foregoing analysis takes a longer term view.

3.6.2 Undertaking response analysis over time

Dillon and Anderson (1990) note two approaches to determining best operating conditions when time is an important element of the decision. One is by use of differential calculus, where the overall problem is treated as a single problem in many variables. The other, known as Dynamic Programming (DP), involves approaching the problem as a series of recursively-related problems, each involving a few variables.

A dynamic optimisation methodology such as DP is useful for strategic and tactical management decisions with the following characteristics:
- the outcomes of decisions in one year, plus the impacts of climatic events, influence the 'state of nature' of a resource or stock in the subsequent year;
- decisions at points of time depend on uncertain events (such as rainfall); and
- the aim is to maximise an objective such as profits while accounting for the biological changes to stocks that occur because of the regular decisions.

3.6.3 Optimal fertiliser applications for wheat fallow rotations

Because of the effects of the Nitrogen cycle and the carryover effects from one year to the next, the tactical farm management decision about fertiliser application for wheat crops has impacts beyond that year (Kennedy et al. 1973, Kennedy 1981a,b, 1986, 1988) . The effect of alternative crop sequences (where pulse crops or pasture legumes might be included) on mineral nitrogen status can also be important in a

strategic sense. In addition the effects of tillage methods and fallow management affect the fertility of the soil.

The effects of crop rotation, tillage method and fertiliser application decisions contribute to dynamic processes, whereby current decisions have impacts on future conditions or states. The stochastic nature of product prices and climatic events are a complicating factor. At the time of making these decisions, the possible outcomes are not known with certainty. Therefore the problem is stochastic and returns should be viewed in an expected value framework.

An optimal strategy for choice of crop rotation and level of fertiliser application involves determining a sequence of decisions so that the expected present value of net returns is maximised, possibly subject to risk exposure constraints. DP is an appropriate method to study a stochastic dynamic problem such as this.

There are a number of examples of analyses that have considered the dynamic and stochastic nature of this question. Taylor (1983) initially used certainty equivalence in determining optimal fertiliser application rates with carryover. Stauber, Burt and Linse (1989) considered the issue of optimal nitrogen fertilisation for seeded grasses in semiarid regions where carryover is significant, and used a stochastic DP approach. Chiao and Gillingham (1989) also developed a DP model to examine the effects on yield and profitability of variability in spread of phosphate fertiliser under carryover conditions and the value of technology which reduces this variability. Segarra et al. (1989) developed a dynamic optimisation model which introduced an intertemporal nitrate-nitrogen residual function to derive and evaluate optimal fertiliser decision rules for irrigated cotton production in Texas. They found that single-year optimisation led to sub-optimal nitrogen applications , which helped explain long-term cotton yield declines in the region. Jomini et al. (1991) modelled fertiliser response in a static and dynamic framework using a LRP function for pearl millet production in Niger. Kim and Hostetler (1991) developed a water-constrained optimal control model of nitrogen fertiliser use, which they used to discuss the effects of a user tax and a government subsidy program as price guided policies for pollution control. Dai et al. (1993) modelled crop decisions, especially for nitrogen fertilisation and soil moisture, for a number of different soil and water conditions in Columbia, Missouri. They incorporated the effects of stochastic factors such as rainfall and temperature on yield response to input changes. Feinerman and Voet (1995) used stochastic DP to evaluate optimal long-run fertilisation and irrigation policies for a perennial crop. In particular they evaluated the economic impact of extra information about the level of leaf nitrogen from leaf tissue analysis. Yadav (1997) used an economic relationship between agricultural production and groundwater pollution to investigate the optimal nitrogen application rates to maintain nitrate contamination in groundwater at a particular standard level over time. He showed that when accounting for groundwater contamination, the best fertilisation rates were less than the profit maximising level.

In this analysis the decision is confined to fertiliser application to a wheat fallow rotation. We abstract from questions of alternative crop rotations, fallow lengths, tillage methods, soil water planting rules or pollution issues. The model can be subsequently expanded to address some of these questions.

3.6.4 Soil Nitrogen dynamics

The DP model must consider the effects of carryover of soil fertility from one crop/year to the next. A critical component of this equation is the nitrogen removed with harvested grain ($GrainN_t$), at time t, which depends on the yield (Y) and protein (PR) level of the crop:

$$GrainN_t = Y_t * PR_t * 1.75. \tag{11}$$

The amount of nitrogen left unused after harvest results from a number of processes which occur during the crop growth period, and which are modelled simultaneously within APSIM. The relationship explaining the change in soil available (or mineral) nitrogen during the cropping period is:

$$N_H = N_t + FertiliserN_t - GrainN_t - StrawN_t - RootsN_t. \\ - Denit / leachN_t + NewMinN_t \tag{12}$$

In this equation, the nitrogen remaining after harvest (N_H) equals nitrogen at sowing (N_t) plus added nitrogen ($FertiliserN_t$), less nitrogen removed with grain ($GrainN_t$), or tied up in straw ($StrawN_t$) and roots ($RootsN_t$), less losses during the crop by denitrification and leaching processes ($Dnit / leachN_t$) plus nitrogen newly mineralised during the crop period ($NewMinN_t$). Estimates of N_H from APSIM are used in the analysis.

The change in mineral nitrogen over the subsequent fallow period was also estimated. This depends on the amount left over after the previous harvest, and the type of rainfall/climate experienced during the fallow. Rather than producing the outcomes for each of the 250 yield/protein figures, a simplified rule was used for fallow losses and gains.

Fallow losses (denitrification and leaching) were assumed to be 10% of N_H. Fallow gains or net mineralisation (FM_t) were assumed to be 17, 27, 30, 33 and 40 units of mineral nitrogen depending on the rainfall during the fallow period following the crop, where 17 units corresponds to a 10 percentile rainfall and 40 units to the 90 percentile rainfall. Note that these numbers are soil specific but are averaged over the crop stubble residue loads. Therefore they do not exactly relate to the residue loads likely to have resulted from specific applied nitrogen decisions in the previous crop.

The expected relationship between soil available nitrogen at subsequent sowing periods is:

$$N_{t+1} = 0.9. N_H + FM_t, \text{ or} \tag{13}$$
$$N_{t+1} = 0.9.(N_t + FertiliserN_t.E - GrainN_t - StrawN_t - RootsN_t \\ - Denit / leachN_t + NewMinN_t) + FM_t$$

Here N_t is part of N_H in (12), and FM_t is the fallow gain, or net mineralisation, which depends on fallow rainfall and comprises the figures (17, .. , 40) above. This relationship represents the change in soil available nitrogen levels between crops, which depends on the fertiliser nitrogen decision in each year, on in-crop and fallow seasonal conditions, and on fallow mineralisation conditions.

3.6.5 The general decision model and recursive equation

In the model used here the problem is to select fertiliser applications out of a set of 25 rates (from 25 to 250 kg/ha in 10 kg intervals). A multi-year planning horizon is used with a decision required annually at planting time. The planning unit is one hectare of crop land, all relevant biological and economic variables are expressed on this basis.

The state variable (N) is the amount of soil mineral nitrogen at the beginning of any sowing stage. The trajectory of the state variable through time depends on the control variable (the amount of fertiliser applied each year, x_t), and on the type of seasons experienced during the planning horizon for both in-crop and fallow periods. The state variable in turn determines the crop response and amount of nitrogen removed at harvest.

The objective function is to maximise the expected present value of net returns from producing wheat over the planning horizon (0, …, T). The recursive relationship for this problem is:

$$V_t(N_t, x_t, C_t, PW_t, PF) = \underset{x}{Max} \left\{ \begin{array}{l} q_t^N \left[Y(N_t, x_t, C_t), PW_t \left(PR_t (N_t, x_t, C_t) \right), PF \right] \\ + \beta E \left[V_{t+1}(N_{t+1}, x_{t+1}, C_{t+1}, PW_{t+1}, PF) \right] \end{array} \right\}, \quad (14)$$

with: $V_{T-1} = q_{T-1}^x(.) + \beta TVAL(N_T)$, $\qquad\qquad\qquad$ (15)

where:

$V_t(\bullet):$ $\quad$ the expected present value of net returns from year t to the end of the planning horizon;

$N_t:$ $\quad$ amount of mineral nitrogen at sowing;

$x_t:$ $\quad$ amount of nitrogen applied;

$PW_t:$ $\quad$ price of wheat;

$C_t:$ $\quad$ type of in-crop rainfall season (very poor, .. , very good);

$PF:$ $\quad$ price of fertiliser;

$Y_t:$ $\quad$ yield of crop;

$PR_t:$ $\quad$ protein content of crop;

$q_t^N(\bullet):$ $\quad$ expected one-period return associated with the current state N;

$\beta:$ $\quad$ discount factor, $1/(1+r)$;

$r:$ $\quad$ discount rate;

$E:$ $\quad$ expectation operator; and

$TVAL(\bullet):$ $\quad$ terminal value of land.

The terminal value in the objective function (equation 15) depends on the fertility in year T. It is assumed that the land market is efficient and pays for soil fertility at fair market value.

3.6.6 The empirical model

In this analysis we use deterministic prices, wheat prices are from May 1999 (Figure 3). The twenty-five seasonal outcomes for yield and quality are assumed to be equally likely.

The model to be solved consists of (14), subject to (12), (13) and (15). The one-period return function is:

$$q_t^N = \overline{PW}_t \left[PR_t(N_t, x_t, C_t) \right] \cdot \overline{Y}_t(N_t, x_t, C_t) - x_t \cdot PF - FVC, \qquad (16)$$

where FVC represents variable costs associated with fallow, and the bars over wheat price and yield variables indicate expectations conditional on the state N_t and the control x_t. The terminal value function ($TVAL$) was approximated as the return vector (16) divided by the discount rate.

4. A synthetic data-set

The APSIM model (McCown et al. 1996) is used in this paper to generate biological outcomes as a basis for discussion and analysis. APSIM is a cropping systems simulation model developed for use as a systems analysis tool for both researchers and growers in the grain cropping regions of North Eastern Australia. The major factors affecting production addressed by this model are climate variability, soil water characteristics, soil nitrogen fertility, variety phenology, planting time and density. APSIM is a relatively complex, daily time step wheat model capable of simulating soil water and nitrogen dynamics over long time spans and under crop rotations with either fixed length fallows or opportunistic sowing rules. For this paper, APSIM has been configured to simulate continuous wheat with summer fallow, where soil fertility and water supply is reset each planting at predetermined levels.

The dataset of outputs from this model is considered to be a good biological representation of responses to added Nitrogen in wheat, and is therefore used in the economic response analysis. It exhibits the characteristics discussed in Section 3, and so is useful in an analysis of the derivation of best operating conditions under each type of departure from the classical response case.

For this exercise, the model output used consists of three variables (wheat yield and protein content are harvested outputs, and residual mineral nitrogen at harvest is another outcome of the process) for each of the 91 simulated years. The soil type chosen is a black vertisol (ie deep cracking) and soil available (mineral) nitrogen is set to 25 kg/ha at planting each year. This enables a response to total soil mineral nitrogen to be estimated. The wheat variety Hartog was sown on 15 July each year using Gunnedah, NSW, climate records.

The model output consists of 91 observations for three variables of interest. These comprise distributions of outcomes, which vary according to seasonal characteristics of rainfall and temperature during crop growth and grain fill periods of the crop, and fallow rainfall conditions.

The model does not deal with phosphorus cycling, so phosphorus supply was assumed to be non-limiting. Similarly, the impacts of frost and disease were assumed to be non-limiting. However, the planting date was chosen when the risk of frost damage would be low.

4.1 Experimental design

The two dominant factors in crop production (provided weed and disease management are adequate) are water and nitrogen supply. While nitrogen supply is manageable through fertiliser addition, variability in water supply (rainfall) is a feature of the northern Australia dryland cropping environment. Rainfall variability in both the in-crop and fallow periods must be considered in this environment as water supply for wheat yield is determined by the level of stored water at planting (the result of fallow rainfall) and in-crop rainfall. Fortunately the amounts of fallow and in-crop rainfall appear to be uncorrellated (Peter Hayman, personal communication).

Simulations of the summer fallow (1900-1990) at Gunnedah have allowed determination of the 10^{th}, 30^{th}, 50^{th}, 70^{th} and 90^{th} percentile values of soil water at planting from frequency distributions. These values correspond to 63mm, 97mm, 124mm, 180mm and 222mm respectively, and have subsequently been referred to as the 'very poor', 'poor', 'average', 'good' and 'very good' fallow rain treatments. Model simulations were conducted over the 1990-1990 period for 50 treatments, ie each year reset at a combination of the 5 starting soil moisture levels by 10 nitrogen supply levels (25, 50, 75, 100, 125,150, 175, 200, 225 and 250 kg nitrogen/ha). The outcomes of the 91 simulated years for each treatment were summarised by determining the 10^{th}, 30^{th}, 50^{th}, 70^{th} and 90^{th} percentile yield, protein and residual nitrogen levels, reflecting the level of in-crop rainfall. The final result is 250 values for each of the output variables resulting from 25 different water supply scenarios (5 fallow season outcomes by 5 in-crop seasons outcomes) and 10 different nitrogen supply scenarios.

4.2 Yield and protein outcomes

The yield responses to total nitrogen (comprising the initial 25 soil available units plus added amounts) are shown graphically for very poor, average and very good fallow rainfall seasons in Figure 5. The corresponding protein responses are also shown there. These figures contain a subset of all results used in this analysis.

In Figure 5, each set of figures (which corresponds to an initial soil moisture level) shows that as extra nitrogen is added the resulting yield generally increases but at a decreasing rate, so that a concave relationship is observed. In general the yield rises at some rate, which depends on the moisture availability, and then reaches a maximum or ceiling level beyond which there is no further gain as extra nitrogen is added. This

level of maximum attainable yield is lower for both drier than wetter fallows and for drier than wetter in-crop rainfall seasons. This type of response is expected from first principles of fertiliser and water requirements in crop growth. As expected from these principles, the yield level reaches its maximum at higher levels of soil fertility as more water is present. That is, the point of maximum yield moves to the right as moisture presence (from whatever source) increases. There are characteristics of the LRP representation in these responses.

For the corresponding protein responses, the general trend is for protein levels to increase as the level of moisture decreases. This is consistent with the crop growth principles (dryness in the finishing stage) discussed above. There are some slight indications of convexities in the shape of the response function. Another feature of the protein response is that as the level of in-crop rainfall declines, protein limits are reached in some cases, this is a ceiling concept similar to the yield case. As moisture conditions decline, especially for average or below fallow rainfall seasons, the protein content of wheat rises and the plateau moves to the left.

In summary, for the fertiliser responses shown in Figure 5 wheat yield responds to both added nitrogen and extra moisture positively, but the protein percentage responses are positive for added nitrogen but negative for extra moisture. As nitrogen is added, yield responses appear to move up and to the right with wetter seasons, and protein responses appear to move up with drier seasons. Thus the probabilities of different seasonal outcomes and the interactions of yield with quality and the payment basis for wheat grain make for a complicated decision by farm managers.

4.3 Soil available nitrogen outcomes

The third outcome from the APSIM results is the amount of soil fertility left after the crop is harvested. This is the mineral nitrogen outcome resulting from crop inputs and the predicted outputs based on weather. This includes both the likely mineral nitrogen available (or unused) after harvest and the net additions to those amounts in the subsequent fallow. Figure 5 also contains a plot of the estimates of unused nitrogen after harvest for very poor, average and very good fallow rainfall seasons.

4.4 Estimated response functions

The APSIM model was used to obtain estimates of wheat yield, protein levels and soil mineral nitrogen after harvest at various levels of fertility. These data were then used to estimate a response function based on the Mitscherlich (1909) equation:

$$Y = \alpha + (\beta - \alpha)\frac{1 - k^{\frac{(N-1).(TotalN-100)}{100}}}{1 - k^{(N-1)}}. \tag{17}$$

The parameter values for α, β and k are given in Table 1 for the different climatic combinations which each provided a different response. The responses describing nitrogen left at harvest (measuring N_H in (12) and (13)) were generally positively linear above a certain level of N.

20

5. Response analyses accounting for the exceptions

In this section the analytical methods available to address the issues in Section 3 using the dataset described in Section 4 are presented, along with the results from the case study analysis.

5.1 Non-smooth response and multiple non-separable outputs

An example of a net benefit curve and marginal rate of return analysis, based on the case study data set, is shown in Figure 6. This applies to immediate responses to added nitrogen, without considering potential carryover effects. For an average in-crop and fallow season, the marginal return analysis indicates that using 100 units of nitrogen provides a marginal rate of return of 97%. Below the level of 100 units, the rate is higher, and beyond that the return is negative.

A plot of adjusted GM against input level could be used as an indication of best operating conditions. Figure 7 shows GM and adjusted GM after requiring a 50% return on invested capital. In the first case the highest profit is realised at 125 units, but with a 50% target rate the highest profit is shown at 100 units of total Nitrogen.

Neither the marginal rate of return nor the profit function approaches overcome the third major problem with classical economic response analysis; the issue of uncertainty in production and price responses and the task of determining best operating conditions likely to be preferred by risk averse decision makers. These results are presented next.

5.2 Accounting for variability in response and price

Figure 5 contains the yield and protein responses to different levels of total nitrogen from APSIM for different in-crop and fallow climate outcomes. There are clear differences in response between in-crop season types, and the divergence for yield increases as higher levels of fertiliser are applied, although this is not the case for protein. Yield increases and protein decreases with improved seasonal conditions.

5.2.1 Profit functions and marginal analysis

When these two sets of responses are combined into a GM using prices like those in Figure 3, the resulting trends for each climate type are as shown in Figure 8 which shows the situation for an average fallow rain season. The distribution of returns widens dramatically as levels of input are used up to about 150 units. Similar divergent patterns are present for other fallow rain patterns. The problem becomes to identify the best input level amongst this variability.

The stochastic problem for the decision maker can also be illustrated using the marginal rate of return analysis. By applying a target rate of return to partial budgets for each climate type (in-crop and fallow seasons) for the yield and protein responses, the best operating conditions for different nitrogen application are shown in Table 2. In that table two sets of farm-gate prices are used, and two target rates of return (25%

and 50%) are analysed. The prices are from December 1997 (with the base (13% protein) Prime Hard price of $230/t fob) and May 1999 (with base Prime Hard price of $153/t fob). The results are for total nitrogen used by the crop. The wheatgrower would need to determine how much available fertiliser is in the soil prior to sowing and then add artificial nitrogen to make up the total amount shown in the table.

A number of observations can be made from the results in Table 2. The first is that there is a range of optimum total fertiliser input levels. Depending on seasonal conditions (both for fallow and in-crop periods) total nitrogen required varies from 75 to 200 units. The second is that the optima do not vary much between the 25 and 50% target rates of return. The marginal analysis in this case seems to give reasonably consistent answers for these target rates. The third is that the optimum rate does not seem to change much for a lower versus a higher price. Thus price variation *per se* does not seem to influence input levels very much. However, it must be noted that these results are developed for particular sets of prices, and we cannot say how general they might be.

5.2.2 Results of Generalised Stochastic Dominance analysis

The distributions of some of the GMs for added nitrogen levels from Figure 8 are presented as CDFs in Figure 9. The results of the analysis using RISKROOT to compare the distributions are shown in Table 3. Based on the premise that the results of the fallow rainfall are observable at the time of crop planting, the stochastic dominance results are presented separately for each fallow season type. The results are presented for an absolute RAC level of $2.4x10^{-6}$, as discussed above. Two sets of results are shown, for GM distributions with a target rate of return of zero and 50%.

For this level of risk aversion a zero rate of return indicates optimal levels of total nitrogen range from 100 to 150 units for very poor to very good fallow rainfall seasons. Thus the optimal level varies according to initial soil moisture conditions. When a 50% target rate of return is specified, the optimal rates are reduced by 25 units in the very poor and good seasons.

5.3 Accounting for the dynamics of nitrogen carryover

The DP results are presented in Table 4 and Figures 10, 11 and 12. With a terminal value of soil nitrogen specified as shown above, variable costs (apart from nitrogen) of $170/ha and a fertiliser price of $0.80/kg (Scott 1998), the model was solved using a Matlab program (The MathWorks Inc. 1998). Initially, the model was solved for the deterministic case over all combinations of in-crop and fallow seasons, and for the higher priced scenario. The problem was solved over an 8-year horizon, this was long enough to obtain a convergent solution for both variables. Then a stochastic model was solved.

5.3.1 Deterministic results

The results are presented in Table 4 as the optimal stock levels of soil fertility and optimal decisions for nitrogen application. An example of the optimal state output for an average fallow and average in-crop rain season is shown in Figure 10. In all cases

results for optimal state and decision converge to a single value from any initial level of soil fertility. The increase in $N*$ in the final year is caused by the final value function and can be ignored in the analysis of longer planning horizons.

The results on Table 4 are for the case where a specific type of fallow season is followed by a specific type of in-crop season for every year of the planning horizon considered. As such, the results are unrealistic since the randomness of climate is ignored. However, the results provide some insights in the process of setting up and solving a more representative DP model.

The DP results can be interpreted as follows. Taking the average in-crop and fallow season case of $N* = 110$ and $x* = 75$ units, we can say that the optimal mineral nitrogen stock prior to sowing in any year is 110 units. The economically optimal amount of fertiliser nitrogen to apply is 75 units, and in this certain world the protein removal in harvested grain plus carryover effects ensure that 110 units are present prior to the next sowing. These results are derived from the combined effects of yield and protein outputs and the nitrogen carryover outcomes shown in Figure 5.

The static marginal results for the same case are given in Table 2, where 125 fertility units in total is the optimal level. Remember that this result is derived from the same APSIM outputs for yield and protein. The interpretation of this figure is that the farm decision maker must establish how much mineral nitrogen is present in the soil prior to sowing and add the extra to make the total to be 125. The marginal analysis is static in the sense that it does not incorporate the carryover effects shown in Figure 5.

In comparing the two results we see that the optimal stock level for the dynamic case is lower. This is expected because of the incorporation of carryover effects. Similar results are reported by Kennedy.

Optimal average nitrogen input levels ($x*$), across in-crop rainfall seasons, were 55, 68, 76, 100 and 101 units applied for very poor, poor, average, good and very good fallow rainfall seasons.

In comparison with the generalised stochastic dominance results in Table 3, the DP results are also generally lower. The stochastic dominance results for an average fallow season indicate a soil fertility level of 125 units. From Table 4 the DP results provide an optimal state of up to 106 units, for any in-crop rainfall season. The stochastic dominance results exclude the carryover effects, so they would be expected to be higher. However, the inclusion of risk aversion would normally lead to a lower optimal rate. The actual outcome would depend on the specific assumptions regarding the level of risk aversion relative to the strength of the carryover effect.

5.3.2 Stochastic results

The stochastic DP model incorporates the expected value framework, so that the response information within any crop season involves weighted probabilities of each in-crop outcome (yield, protein and carryover). In the return function (16), expected profit is now expressed as a probability-weighted sum of price and yield less variable costs. In practice, the DP model is solved by utilising a probability transition matrix for fallow rainfall (FR) and in-crop rainfall (ICR) seasons. The stochastic DP model

solves over random fallow rainfalls (initial soil moisture conditions) by taking expectations over each in-crop rainfall outcome range.

The output from this model, for both the state and control variables, comprises a three dimensional matrix of dimensions (N, FR, ICR). A characteristic of this model is that, although over the planning horizon both ICR and FR are random variables, within each decision period t, FR_t is known at planting time, and this information can be used when making fertiliser application decisions. In other words, the fertiliser decision can be improved upon by using information on the current FR outcome. When retrieving the optimal path, a given sequence of ICR and FR values is assumed, and the decision-rule matrix produced by the DP model solution is used to obtain the particular optimal solution. During this process each FR_t is given its actual value and the optimal value of x_t is determined based only on possible ICR outcomes.

To study the stochastic results, a Monte Carlo simulation of 1000 different vectors of random FR and ICR combinations was undertaken. The plots in Figure 11 (solid lines) were derived by averaging the optimal results from these 1000 possible climatic sequences. Note that the assumption of risk neutrality is implicit in the derivation of these solutions. To represent other risk attitudes the objective function must be modified to maximise utility rather than profit.

The same 1000 climatic sequences were applied to the DP decision rule for the deterministic case with average ICR and FR values. The results are shown in Figure 11 (dotted lines). The stochastic dynamic programming yields slightly higher optimal equilibrium state and control values than the deterministic model. This is plausible because the inclusion of seasonal response functions and downside risk would make it reasonable to maintain a higher fertility level to account for the possibility of good seasons. The nitrogen carryover process reinforces this decision.

Differences in optimal state and control values ($N^* = 108$ kg and $x^* = 77$ kg in the deterministic case compared to $N^* = 122$ kg and $x^* = 83$ kg in the stochastic case) result in a very small difference in average returns (compare solid and dotted lines for q^* in Figure 11). The discounted annual returns are initially over $450, and decline over the 15 years. The initial return is well within the range of other estimates for the assumed prices and yield outcome (Scott 1998). The return declines over time due to the effects of discounting.

The small differences in returns suggest that the use of a certainty equivalent is a good approximation to the solution of this problem, which implies that it may not be necessary to explicitly account for stochasticity of the environment in this particular problem. This is good news because it means that the high cost of solving the stochastic model can be avoided. This result can be explained through the carryover process, which ensures that a 'bad' fertiliser decision in one year still provides benefits in future years. It is optimal to maintain soil fertility at a level that will allow the firm to take advantage of good years.

The distribution of the present value of profits for a 15-year period is shown in Figure 12. The 50[th] percentile value is just under $6000. The stochastic returns lie marginally (but consistently) to the right of the deterministic returns, however the

distributions are very close. The advantage of using a stochastic solution approach may become relevant at higher income target levels. The probability that the present value of returns will be above $6,000/ha, for example, is 0.28 in the deterministic solution and increases to 0.33 in the stochastic solution. Whether this five percent difference in the probability of obtaining higher profits is enough to warrant the extra cost of using a stochastic model is an open question. The present discounted value is an indication of the market value of the fixed resource (land) when used optimally to grow short fallow wheat. Extending the time frame beyond 15 years would shift the distribution slightly to the right.

5.4 Comparison with other studies

Turpin et al. (1998) and Hayman and Turpin (1998) have analysed nitrogen fertiliser decisions for the same location when considering stored soil water and climate forecasts, and paddock history and soil tests, respectively. They used APSIM to investigate the effects of using fixed versus flexible fertiliser strategies. With respect to prior information on stored soil water and climate forecasts, Hayman and Turpin suggested that farmers aiming to apply between 85 and 100 units of nitrogen are on a relatively flat part of the economic response curve, and that by selecting a fixed application rate in this range they would not sacrifice too much potential gain in good seasons nor lose too much in poor seasons. Once farmers apply these amounts, the flatness of the curve means that adjusting nitrogen rates has a small impact on GM.

While it can be difficult to compare between analyses which are conducted with different aims, a couple of observations can be made about these other results compared to the present study. First, Turpin et al.'s results for nitrogen at planting and rate of nitrogen fertiliser appear to be slightly higher than the optimal state and application decision figures in Table 4. Second, Hayman and Turpin's range of 85 to 100 units of applied nitrogen is also slightly higher than the range of x^* in Table 4.

The reasons for any differences are unclear because of the methods used and factors incorporated into each analysis. There is scope for further cooperative work here.

6. Discussion

In this paper the economic theory and practice of response analysis are reviewed and applied to a synthetic crop response data set. A range of methods that can be applied to real-world questions are illustrated. Our results are consistent with other analyses and also reasonable in terms of *a priori* expectations.

The DP model of wheat production we develop accounts for grain quality (protein content) and nitrogen carryover between time periods. Results suggest that it may not be necessary to account for the stochastic production environment in this type of problem, because of the forgiving nature of nitrogen application outcomes in the presence of carryover. The effect of stochastic prices was not explored and remains an important topic for future research.

In closing, we point out that the applied agricultural R&D industry does not appear to utilise marginalist economic thinking very widely. A more widespread usage of some

of the methods discussed here in the agricultural research, development and extension community, where technology and adoption by groups of industries is the primary objective, may be desirable. However, it may be some time before the more complex stochastic and dynamic models developed here will be adopted.

References

Anderson, J.R. (1967), 'Economic interpretation of fertiliser response data', *Review of Marketing and Agricultural Economics* 35(1), 43-57.

Anderson, J.R. (1975), 'One more or less cheer for optimality', *Journal of the Australian Institute of Agricultural Science*, 41(3): 195-7.

Anderson, Jock R., John L. Dillon and J. Brian Hardaker (1977), *Agricultural Decision Analysis*, The Iowa State University Press, Ames Iowa.

ABARE (1999), *Australian Farm Surveys Report 1999*, Canberra.

Bar-Shira, Z., Just, R.E. and Zilberman, D, (1997), 'Estimation of farmers' risk attitude: an econometric approach' *Agricultural Economics* 17: 211-222.

Bardsley, P and Harris, M (1987), 'An approach to the econometric estimation of attitudes to risk in agriculture', *Australian Journal of Agricultural Economics*, 31(2), 112-26.

Chiao, Yen-Shong and Gillingham, Allan (1989), 'The value of stabilising fertiliser under carry-over conditions', *American Journal of Agricultural Economics*, 352-362.

CIMMYT (1988), *From Agronomic Data to Farmer Recommendations: An Economics Training Manual*, Completely revised edition, Mexico D.F.

Dai, Qi, Fletcher, Jerald J. and Lee, John G. (1993), 'Incorporating stochastic variables in crop response models: implications for fertilisation decisions', *American Journal of Agricultural Economics* 75, 377-386.

Dalal, R, and Mayer, R. (1986), 'Long term trends in fertility of soils under continuous cultivation and cereal cropping in southern Queensland. I. Overall Changes in Soil Properties and Trends in Winter Cereal Yields', *Australian Journal of Soil Research* 24: 265-279.

Dillon, John L. and Jock R. Anderson (1990), *The Analysis of Response in Crop and Livestock Production*, Third Edition, Pergamon Press,

Feinerman, Eli and Voet, Hillary (1995), 'Dynamic optimisation of nitrogen fertilisation of citrus and the value of information from leaf tissue analysis', *European Review of Agricultural Economics* 22, 103-118.

France, J. and J.H.M. Thornley (1984), *Mathematical Models in Agriculture: A Quantitative Approach to Problems in Agriculture and Related Sciences* , Butterworths.

Fraser, Rob (1998), 'Seasonal variability and a farmer's supply response to protein premiums and discounts', *The Australian Journal of Agricultural and Resource Economics,* 42(1): 25-33.

Hammond, J.S. (1974), 'Simplifying the Choice Between Uncertain Prospects Where Preference Is Nonlinear' *Management Science* 20, no. 7: 1047-72.

Hayman, Peter and Claire Alston (1999), 'A survey of farmer practices and attitudes to nitrogen management in the northern New South Wales grain belt' *Australian Journal of Experimental Agriculture* 39: 51-63.

Hayman, P.T. and J.E. Turpin (1998), 'Nitrogen fertiliser decisions for wheat on the Liverpool Plains, NSW. II. Should farmers consider stored soil water and climate forecasts?', Proceedings of the 9th Australian Agronomy Conference, Wagga Wagga.

Hayman, Peter and James de Vries (1995), 'Managing Crop Sequences: A guide to your options in North West NSW', NSW Agriculture, Orange.

Jauregui, M.A. and G.E. Sain (1992), *Continuous Economic Analysis of Crop Response to Fertiliser in On-Farm Research*, CIMMYT Economics Paper No. 3, Mexico, D.F.: CIMMYT.

Jomini, Patrick A., Deuson, Robert R., Lowenberg-DeBoer, J. and Bationo, Andre (1991), 'Modelling stochastic crop response to fertilisation when carry-over matters', *Agricultural Economics* 6, 97-113.

Just, Richard E. and Quinn Weninger (1999), 'Are Crop Yields Normally Distributed?' *American Journal of Agricultural Economics* 81: 287-304.

Kennedy, John O.S. (1981a), 'Applications of Dynamic Programming to Agriculture, Forestry and Fisheries: Review and Prognosis', *Review of Marketing and Agriculture and Economics* 49(3), 141-173.

——— (1981b), 'An Alternative Method for Deriving Optimal Fertiliser Rates: Comment and Extension', *Review of Marketing and Agriculture and Economics* 49(3), 203-209.

——— (1986), 'Rules for Optimal Fertiliser Carryover: An Alternative Explanation', *Review of Marketing and Agriculture and Economics* 54(2), 3-10.

——— (1988), 'Principles of Dynamic Optimisation in Resource Management', *Agricultural Economics* 2, 57-72.

Kennedy, J.O.S., Whan, I.F., Jackson, R. and Dillon, J.L. (1973), 'Optimal fertiliser carryover and crop recycling policies for a tropical grain crop', *Australian Journal*

of Agricultural Economics 17(2), 104-113.

Kim, C.S. and Hostetler, John E. (1991), 'Policy implications of a water quality-constrained dynamic model of nitrogen fertiliser use', *Canadian Journal of Agricultural Economics* 39, 781-791.

Lanzer, E.A. and Paris, Q. (1981), 'A new analytical framework for the fertilisation problem', *American Journal of Agricultural Economics* 63, 93-103.

Makeham, J.P. and Malcolm, L.R. (1981), *The Farming Game*, Gill Publications, Armidale.

Martin, Bob, Harry Marcellus, Andrew Verrell and David Herridge (eds) (1996), 'Grower's Workbook, Nitrogen in '96: Decisions about nitrogen fertiliser for high yield and quality wheat', Grains Research and Development Corporation and NSW Agriculture, Orange, February.

McCarl, Bruce A. (1988), 'Preference Among Risky Prospects Under Constant Risk Aversion' *Southern Journal of Agricultural Economics*: 25-33.

——— (1990) "Generalised Stochastic Dominance: An Empirical Examination." *Southern Journal of Agricultural Economics*: 49-55.

McCown, R.L., G.L. Hammer, J.N.G. Hargreaves, D.P. Holzworth and D.M. Freebairn (1996), 'APSIM: a Novel Software System for Model Development, Model Testing and Simulation in Agricultural Systems Research', *Agricultural Systems*, 50: 255-271.

Meyer, J. (1977), 'Choice Among Distributions' *Journal of Economic Theory* 14: 326-36.

Mitscherlich, E.A. (1909), 'Das Gesetz des Minimums un das Gesetz des abnehmenden Bodenertrages', *Landwirtschaftliche Jahrbucher* 38: 537-552.

Parton, K.A. and P.S. Carberry (1995), "Stochastic Dominance and Mean-Standard Deviation Analysis: Some Critical Issues*" Australian Journal of Agricultural Research* 46: 1487-1491.

Patten, Louise H., J. Brian Hardaker, and David J. Pannell (1988), 'Utility-Efficient Programming for Whole-Farm Planning' *Australian Journal of Agricultural Economics* 32, no. 2 and 3: 88-97.

Pratt, J.W. (1964), 'Risk Aversion in the Small and the Large' *Econometrica* 32, no. 1-2: 122-36.

Scott, F. (1998), *Farm Budget Handbook 1997 Northern NSW - Winter Crops*, NSW Agriculture, Tamworth.

Segarra, Eduardo, Ethridge, Don E., Deusson, Curtis R. and Onken, Arthur B. (1989), 'Nitrogen carry-over impacts in irrigated cotton production, Southern High Plains of Texas', *Western Journal of Agricultural Economics* 14(2), 300-309.

Stauber, M.S., Burt, O.R. and Linse, F. (1989), 'An economic evaluation of nitrogen fertilisation of grasses when carryover is significant', *American Journal of Agricultural Economics* 57, 463-471.

Strong, W.M. (1981), 'Nitrogen requirements of irrigated wheat on the Darling Downs', *Australian Journal of Experimental Agriculture and Animal Husbandry*, 21: 424-31.

Taylor, C.R. (1983), 'Certainty equivalence for determination of optimal fertilisation rates with carryover', *Western Journal of Agricultural Economics* 8, 64-67.

The MathWorks Inc. (1998), *Matlab: The Language of Technical Computing*, Version 5.2.0.3084.

Turpin, J.E., P.T. Hayman, H. Marcellus and D.F. Freebairn (1998), 'Nitrogen fertiliser decisions for wheat on the Liverpool Plains, NSW. I. Should farmers consider paddock history and soil tests?', Proceedings of the 9[th] Australian Agronomy Conference, Wagga Wagga.

Waggoner, P.E. and Norvell, W.A. (1979), 'Fitting the Law of the Minimum to Fertiliser applications and Crop Yields', *Agronomy Journal* 71: 352-354.

Yadav, Satya N. (1997), 'Dynamic optimisation of nitrogen use when groundwater contamination is internalised at the standard in the long run', *American Journal of Agricultural Economics* 79, 931-945.

Table 1
Smooth functions fitted to APSIM results

Fallow rain	In-crop rain	Yield response (a)			Protein response (a)			Nitrogen left (b)	
		α	β	k	α	β	k	Intercept	Slope
Very poor	Very poor	1.007	(c)	(c)	15.185	15.323	0.417	-54.56	1.010
	Poor	1.717	(c)	(c)	13.023	13.427	0.563	-68.05	0.995
	Average	2.202	2.285	0.575	11.348	12.730	0.791	-66.76	0.930
	Good	2.775	2.777	0.251	10.735	11.481	0.713	-81.13	0.964
	Very good	3.508	3.780	0.643	9.880	11.130	0.835	-94.63	0.914
Poor	Very poor	1.466	(c)	(c)	14.256	14.358	0.403	-63.55	0.989
	Poor	2.279	2.289	0.346	11.891	12.737	0.690	-76.13	0.968
	Average	2.776	2.815	0.471	10.773	11.757	0.750	-76.85	0.935
	Good	3.249	3.414	0.601	10.154	11.226	0.797	-92.06	0.954
	Very good	3.756	4.272	0.717	9.615	11.159	0.881	-91.17	0.846
Average	Very poor	1.996	2.150	0.245	12.844	13.543	0.634	-77.26	0.981
	Poor	2.793	2.844	0.521	10.814	12.093	0.797	-74.51	0.919
	Average	3.152	3.280	0.596	10.271	11.588	0.825	-82.63	0.928
	Good	3.537	3.816	0.663	9.788	11.108	0.854	-87.96	0.906
	Very good	3.939	4.682	0.770	9.314	10.978	0.918	-82.89	0.780
Good	Very poor	3.204	3.394	0.644	10.524	11.915	0.822	-80.35	0.905
	Poor	3.489	3.844	0.712	9.920	11.366	0.867	-80.68	0.877
	Average	3.698	4.145	0.727	9.613	10.977	0.872	-85.83	0.881
	Good	3.942	4.684	0.781	9.372	10.982	0.917	-80.33	0.790
	Very good	4.257	5.435	0.826	9.078	10.573	0.941	-75.92	0.694
Very good	Very poor	3.380	3.672	0.667	10.524	11.915	0.822	-84.71	0.928
	Poor	3.586	4.088	0.734	9.201	11.366	0.867	-81.73	0.885
	Average	3.759	4.390	0.757	9.613	10.977	0.872	-81.51	0.859
	Good	3.954	4.866	0.798	9.372	10.982	0.917	-79.05	0.796
	Very good	4.242	5.645	0.842	9.078	10.573	0.941	-75.59	0.709

(a) Mitscherlich curve.
(b) Linear form.
(c) Linear response, intercept parameter only, zero slope.
Claire Alston and Steve Harden of NSW Agriculture estimated these parameters.

Table 2
Optimal total fertiliser levels from marginal analysis
Target rate of return: 25 and 50%
Prices: December 1997 and May 1999
Different in-crop and fallow seasons

Fallow	In-crop rain				
rain	Very poor	Poor	Average	Good	Very Good
	Prices	*12/97 (higher)*	*Target rate*	*25%*	
Very poor	75	75	75	100	125
Poor	75	75	100	125	150
Average	75	100	125	125	175
Good	125	125	125	175	200
Very good	125	150	150	150	200
	Prices	*12/97 (higher)*	*Target rate*	*50%*	
Very poor	75	75	75	100	125
Poor	75	75	100	100	150
Average	75	100	125	125	150
Good	125	125	125	150	200
Very good	125	150	150	150	200
	Prices	*5/99 (lower)*	*Target rate*	*25%*	
Very poor	75	75	75	100	125
Poor	75	75	100	100	125
Average	75	100	100	125	150
Good	125	125	125	150	175
Very good	125	125	150	150	200
	Prices	*5/99 (lower)*	*Target rate*	*50%*	
Very poor	75	75	75	100	125
Poor	75	75	100	100	125
Average	75	75	100	100	150
Good	100	100	125	150	175
Very good	125	125	125	150	200

Units are kg nitrogen/ha wheat sown

Table 3
Optimal total fertiliser levels
Generalised Stochastic Dominance analysis
Distributions of profits for different season types
Zero and 50% target rate of return
Risk Aversion Coefficient of $2.4x10^{-6}$

Zero required rate of return		50% required rate of return	
Fallow Rain Category	*Dominant set Total N required*	*Fallow Rain Category*	*Dominant set Total N required*
Very poor	100 units	Very poor	75 units
Poor	100 units	Poor	100 units
Average	125 units	Average	125 units
Good	150 units	Good	125 units
Very good	150 units	Very good	150 units

Table 4
Deterministic Dynamic Programming
Convergent results for fallow and in-crop season categories
Optimal fertiliser input levels and soil fertility levels
Units are soil available Nitrogen kg/ha

Fallow rain	In-crop rain					
	Very poor	Poor	Average	Good	Very good	Mean
	Optimal state (N*)					
Very poor	67.3	62.8	108.0	56.8	96.8	78.3
Poor	48.5	87.8	99.7	88.8	111.6	87.3
Average	62.0	74.6	110.1	101.7	106.3	91.0
Good	140.9	135.0	115.3	110.6	114.7	123.3
Very good	141.7	135.7	122.7	134.4	118.3	130.5
	Optimal control (x*)					
Very poor	27.7	42.2	57.0	58.2	88.2	54.7
Poor	36.5	57.2	65.3	76.2	103.4	67.7
Average	53.0	60.4	74.9	83.3	108.7	76.0
Good	84.1	90.0	89.7	104.4	130.3	99.7
Very good	83.3	89.3	92.3	110.6	126.7	100.5