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.*

ELSEVIER

Agricultural Economics 16 (1997) 139–152

AGRICULTURAL
ECONOMICS

Sensitivity analysis of normative economic models: theoretical framework and practical strategies

David J. Pannell

Agricultural and Resource Economics, University of Western Australia, Nedlands, W.A. 6907, Australia

Accepted 24 October 1996

Abstract

The parameter values and assumptions of any economic model are subject to change and error. Sensitivity analysis (SA), broadly defined, is the investigation of these potential changes and errors and their impacts on conclusions to be drawn from the model. There is a very large literature on procedures and techniques for SA, but it includes almost nothing from economists. This paper is a selective review and overview of theoretical and methodological issues in SA. There are many possible uses of SA, described here within the categories of decision support, communication, increased understanding or quantification of the system, and model development. The paper focuses somewhat on decision support. It is argued that even the simplest approaches to SA can be theoretically respectable in decision support if they are applied and interpreted in a way consistent with Bayesian decision theory. This is not to say that SA results should be formally subjected to a Bayesian decision analysis, but that an understanding of Bayesian probability revision will help the modeller plan and interpret an SA. Many different approaches to SA are described, varying in the experimental design used and in the way results are processed. Possible overall strategies for conducting SA are suggested. It is proposed that when using SA for decision support, it can be very helpful to attempt to identify which of the following forms of recommendation is most appropriate: (a) do X , (b) do either X or Y depending on the circumstances, (c) do either X or Y , whichever you like, (d) if in doubt, do X . A system for reporting and discussing SA results is recommended. © 1997 Elsevier Science B.V.

1. Introduction

The parameter values and assumptions of any economic model are subject to change and error. Sensitivity analysis (SA), broadly defined, is the investigation of these potential changes and errors and their impacts on conclusions to be drawn from the model (e.g. Baird, 1989). SA can be easy to do, easy to understand, and easy to communicate. It is possibly the most useful and most widely used technique available to applied economists (including agricultural economists). The importance and usefulness of SA is widely recognised:

“A methodology for conducting a [sensitivity] analy-

sis...is a well established requirement of any scientific discipline. A sensitivity and stability analysis should be an integral part of any solution methodology. The status of a solution cannot be understood without such information. This has been well recognised since the inception of scientific inquiry and has been explicitly addressed from the beginning of mathematics” (Fiacco, 1983, p. 3.)

There is a very large and diverse literature on SA, including a number of reviews (e.g. Tzafestas et al., 1988; Rios Insua, 1990; Sobieszczanski-Sobieski, 1990; Eschenbach and Gimpel, 1990; Lomas and Eppel, 1992; Hamby, 1994; Clemson et al., 1995). However, the existing literature is limited in a num-

ber of respects. Most of what has been written about sensitivity analysis has taken a very narrow view of what it is and what it can be useful for. A large proportion of the literature is highly mathematical and rather theoretical in nature. Even those papers with a focus on applied methodology have tended to concentrate on systems and procedures which are relatively time consuming and complex to implement. There has been almost no discussion of procedures and methodological issues for simple approaches to sensitivity analysis. (Eschenbach and McKeague, 1989, is a rare exception.) This is remarkable, considering the usefulness and extremely wide usage of simple approaches.

The other area of notable neglect is the entire discipline of economics. Despite countless applications of SA in the various areas of applied economics (e.g. Dungan and Wilson, 1991; Nordblom et al., 1994), there has been hardly any discussion of methodological issues for SA of economic models. Recent exceptions include Canova (1995), Eschenbach and McKeague (1989), Eschenbach and Gimpel (1990), and Harrison and Vinod (1992). Agricultural economists appear to have contributed nothing to the literature about sensitivity analysis.

This paper is an attempt to redress some of these areas of neglect. The focus is on 'normative' usage of models to develop recommendations for decision makers, such as managers. Many techniques and procedures will be discussed, ranging from simple to complex. While it is acknowledged that some of the complex procedures which have been proposed are potentially of high value, the primary objective of this paper is to provide guidance and advice to improve the rigour and value of relatively simple approaches. It will be argued that even the simplest approaches to SA can be theoretically respectable in decision support if they are applied and interpreted in a way consistent with Bayesian decision theory. The paper is relevant to both optimisation and simulation models used for decision support, although there is a greater emphasis on optimisation models in the discussion.

2. Uses of sensitivity analysis

There is a very wide range of uses to which sensitivity analysis is put. An incomplete list is given

in Table 1. The uses are grouped into four main categories: decision making or development of recommendations for decision makers, communication, increased understanding or quantification of the system, and model development. While all these uses are potentially important, the primary focus of this paper is on making decisions or recommendations.

In all models, parameters are more-or-less uncertain. The modeller is likely to be unsure of their current values and to be even more uncertain about their future values. This applies to things such as prices, costs, productivity, and technology. Uncertainty is one of the primary reasons why sensitivity analysis is helpful in making decisions or recommen-

Table 1
Uses of sensitivity analysis

1	Decision making or development of recommendations for decision makers
1.1	Testing the robustness of an optimal solution
1.2	Identifying critical values, thresholds or break-even values where the optimal strategy changes
1.3	Identifying sensitive or important variables
1.4	Investigating sub-optimal solutions
1.5	Developing flexible recommendations which depend on circumstances
1.6	Comparing the values of simple and complex decision strategies
1.7	Assessing the 'riskiness' of a strategy or scenario
2	Communication
2.1	Making recommendations more credible, understandable, compelling, or persuasive
2.2	Allowing decision makers to select assumptions
2.3	Conveying lack of commitment to any single strategy
3	Increased understanding or quantification of the system
3.1	Estimating relationships between input and output variables
3.2	Understanding relationships between input and output variables
3.3	Developing hypotheses for testing
4	Model development
4.1	Testing the model for validity or accuracy
4.2	Searching for errors in the model
4.3	Simplifying the model
4.4	Calibrating the model
4.5	Coping with poor or missing data
4.6	Prioritising acquisition of information

dations. If parameters are uncertain, sensitivity analysis can give information such as:

1. how robust the optimal solution is in the face of different parameter values (use 1.1 from Table 1),
2. under what circumstances the optimal solution would change (uses 1.2, 1.3, 1.5),
3. how the optimal solution changes in different circumstances (use 3.1),
4. how much worse off the decision makers would be if they ignored the changed circumstances and stayed with the original optimal strategy or some other strategy (uses 1.4, 1.6).

This information is extremely valuable in making a decision or recommendation. If the optimal strategy is *robust* (insensitive to changes in parameters), this allows confidence in implementing or recommending it. On the other hand if it is not robust, sensitivity analysis can be used to indicate how important it is to make the changes to management suggested by the changing optimal solution. Perhaps the base-case solution is only slightly sub-optimal in the plausible range of circumstances, so that it is reasonable to adopt it anyway. Even if the levels of variables in the optimal solution are changed dramatically by a higher or lower parameter value, one should examine the difference in profit (or another relevant objective) between these solutions and the base-case solution. If the objective is hardly affected by these changes in management, a decision maker may be willing to bear the small cost of not altering the strategy for the sake of simplicity.

If the base-case solution is not always acceptable, maybe there is another strategy which is not optimal in the original model but which performs well across the relevant range of circumstances. If there is no single strategy which performs well in all circumstances, SA identifies different strategies for different circumstances and the circumstances (the sets of parameter values) in which the strategy should be changed.

Even if there is no uncertainty about the parameter values, it may be completely certain that they will change in particular ways in different times or places. In a similar way to that outlined above, sensitivity analysis can be used to test whether a simple decision strategy is adequate or whether a complex conditional strategy is worth the trouble.

SA can be used to assess the 'riskiness' of a

strategy or scenario (use 1.7). By observing the range of objective function values for different strategies in different circumstances, the extent of the difference in riskiness can be estimated and subjectively factored into the decision. It is also possible to explicitly represent the trade-off between risk and benefit within the model.

3. Theoretical framework for using sensitivity analysis for decision making

In this discussion, a decision variable is a variable over which the decision maker¹ has control and wishes to select a level, whereas a strategy refers to a set of values for all the decision variables of a model. An optimal strategy is the strategy which maximises the value of the decision maker's objective function (e.g. profit, social welfare, expected utility). It is assumed that the modeller has subjective beliefs (internal beliefs, hunches, or guesses) about the performance of different strategies and about what is the objective of the decision maker who will use the information generated by the model. The modeller's subjective beliefs are influenced by the model but also by other factors; these beliefs may or may not be close to the objective truth.

Bayesian decision theory provides two tools which are helpful in the use of SA for decision support: (a) decision theory provides a framework for comparing strategies under risk or uncertainty and (b) Bayes' rule provides a rigorous and consistent method for revising probability distributions of uncertain output variables after new information is obtained from the SA. Together these elements provide a tool for rational and consistent adjustments to strategies and decisions as new information is obtained. SA is a process of creating new information about alternative strategies. When viewed in this light, the relevance of Bayesian decision theory to SA is obvious.

¹ Although the text refers to a single modeller and a single decision maker, SA is also very useful in cases where multiple modellers and decision makers are involved. For example it allows different perceptions about appropriate parameter values to be tested, and reveals whether the differences matter.

Conceptually, the process of conducting an SA to choose an optimal strategy can proceed as follows. Following an initial run with a ‘base-case’ model which incorporates ‘best-bet’ values of parameters, a belief about the optimal strategy can be formed. This belief is based on the modeller’s perceptions of the probability distributions of profit (or another measure of benefit or welfare) for the preferred strategy and other strategies. Based on decision theory, the initial optimal strategy is the one which maximises the expected value of the objective function, given the initial or ‘prior’ set of subjectively perceived probability distributions of profit for different strategies. These prior distributions could also be used to make statements about the modeller’s level of confidence that the initial strategy is optimal.

Following a sensitivity analysis based on one or more of the techniques outlined later, the modeller employs Bayes’ rule to revise the subjectively perceived probability distributions of profit for different strategies (resulting in a set of ‘posterior’ distributions). Depending on how the perceptions change, the optimal strategy may or may not be altered. The posterior distributions are likely to be less uncertain (although not necessarily less risky),² due to the information obtained from the SA, so the modeller can make improved statements about his or her confidence in the strategy.

For this general view of the process to be valuable, it is not necessary for the modeller to literally use Bayes’ rule or a formal decision theory framework. In my opinion, merely conceptualising the process in this way will probably improve the rigour and consistency of the SA. Even if the modeller operates subjectively as a Bayesian decision theorist, it may be that an unstructured ‘what if?’ approach to the SA is adequate for some studies. On the other hand, the modeller may be encouraged to adopt a structured, explicitly probabilistic approach based on decision theory.

² ‘Uncertainty’ refers to ignorance about the true value or probability distribution of a variable. ‘Risk’ refers to the fact that a variable may be intrinsically stochastic, describable by a probability distribution but not by a particular value. Even if there is no uncertainty about a variable (i.e. we have perfect knowledge about its probability distribution) there is likely to still be risk.

One potential conceptual difficulty with the framework arises when this type of SA is conducted with an optimisation model. A perceived benefit of SA is that it conveniently allows assessment of the consequences of parameter uncertainty, even with a deterministic model. However SA with a deterministic optimisation model most commonly generates only a single optimal result for each combination of parameter values being tested. If, as is normal, the value of the uncertain parameter will not be definitely known until after the strategy is fixed in place, there is in fact a range of possible profit outcomes (a probability distribution of outcomes) for each possible strategy. Thus if a standard SA approach is used to investigate parameter uncertainty in a deterministic optimisation model, the resulting output will not be easy to relate to the Bayesian decision theory framework; it provides only a sub-set of the relevant information. Note that this problem is unlikely to arise if a simulation model is used, since the tendency with a simulation model is to generate a full set of SA results for each strategy under consideration, providing more information about the probability distribution of outcomes for that strategy.

There are three possible responses to this difficulty with optimisation models:

1. Deal with the parameter uncertainty by explicitly representing it within a stochastic model, rather than by using SA with a deterministic model.
2. Constrain the optimisation model to a particular strategy and generate solutions for that strategy for each combination of parameter values. This provides the probability distribution of outcomes for that strategy. Repeat the process for each strategy of interest. In this approach, the model is really being used for simulation rather than optimisation. However the optimisation capacity is still useful for helping select which strategies to simulate.
3. Using subjective judgement and mindful of the correct decision theory approach, estimate the posterior distributions based only on the single optimal result for each scenario. While the quality of posterior distributions obtained in this way is likely to be somewhat lower than those obtained by Approaches 1 or 2, this approach is computationally much easier. In practice, a set of single SA results from an optimisation model could still

be very useful if considered within the type of conceptual framework outlined earlier. An awareness of the inconsistency between the SA results and the Bayesian decision theory framework should at least help the modeller interpret the significance and implications of the results.

4. Approaches to sensitivity analysis

In principle, sensitivity analysis is a simple idea: change the model and observe its behaviour. In practice there are many different possible ways to go about changing and observing the model. The section covers what to vary, what to observe and the experimental design of the SA.

4.1. What to vary

One might choose to vary any or all of the following:

1. the contribution of an activity to the objective,
2. the objective (e.g. minimise risk of failure instead of maximising profit),
3. a constraint limit (e.g. the maximum availability of a resource),
4. the number of constraints (e.g. add or remove a constraint designed to express personal preferences of the decision maker for or against a particular activity),
5. the number of activities (e.g. add or remove an activity), or
6. technical parameters.

Commonly, the approach is to vary the value of a numerical parameter through several levels. In other cases there is uncertainty about a situation with only two possible outcomes; either a certain situation will occur or it will not. Examples include:

- What if the government legislates to ban a particular technology for environmental reasons?
- In a shortest route problem, what if a new free-way were built between two major centres?
- What if a new input or ingredient with unique properties becomes available?

Often this type of question requires some structural changes to the model. Once these changes are made, output from the revised model can be compared with the original solution, or the revised model

can be used in a sensitivity analysis of uncertain parameters to investigate wider implications of the change.

4.2. What to observe

Whichever items the modeller chooses to vary, there are many different aspects of a model output to which attention might be paid:

1. the value of the objective function for the optimal strategy,
2. the value of the objective function for sub-optimal strategies (e.g. strategies which are optimal for other scenarios, or particular strategies suggested by the decision maker),
3. the difference in objective function values between two strategies (e.g. between the optimal strategy and a particular strategy suggested by the decision maker),
4. the values of decision variables,
5. in an optimisation model, the values of shadow costs, constraint slacks or shadow prices, or
6. the rankings of decision variables, shadow costs, etc.

4.3. Experimental design

The experimental design is the combinations of parameters which will be varied and the levels at which they will be set. The modeller must decide whether to vary parameters one at a time, leaving all others at standard or base values, or whether to examine combinations of changes. An important issue in this decision is the relative likelihood of combinations of changes. If two parameters tend to be positively correlated (e.g. the prices of two similar outputs) the possibility that they will both take on relatively high values at the same time is worth considering. Conversely if two parameters are negatively correlated, the modeller should examine high values of one in combination with low values of the other. If there is no systematic relationship between parameters, it may be reasonable to ignore the low risk that they will both differ substantially from their base values at the same time, especially if they are not expected to vary widely.

In selecting the parameter levels which will be used in the sensitivity analysis, a common and nor-

mally adequate approach is to specify values in advance, usually with equal sized intervals between the levels (e.g. Nordblom et al., 1994). The levels selected for each parameter should encompass the range of possible outcomes for that variable, or at least the ‘reasonably likely’ range. What constitutes ‘reasonably likely’ is a subjective choice of the modeller, but one possible approach is to select the maximum and minimum levels such that the probability of an actual value being outside the selected range is 10%.

If combinations of changes to two or more parameters are being analysed, a potential approach is to use a ‘complete factorial’ experimental design, in which the model is solved for all possible combinations of the parameters. While this provides a wealth of information, if there are a number of parameters to analyse, the number of model solutions which must be obtained can be enormous. To conduct a complete factorial sensitivity analysis for eight parameters each with five levels would require 390 625 solutions. If these take 1 min each to process, the task would take 9 months, after which the volume of output created would be too large to be used effectively. In practice one must compromise by reducing the number of variables and/or the number of levels which are included in the complete factorial. Preliminary sensitivity analyses on individual parameters are helpful in deciding which are the most important parameters for inclusion in a complete factorial experiment. (See later comments on ‘screening’.)

Alternatively one may reduce the number of model solutions required by adopting an incomplete design with only a sub-set of the possible combinations included. Possibilities include central composite designs (e.g. Hall and Menz, 1985), Taguchi methods (e.g. Clemson et al., 1995), or some system of random sampling or ‘Monte Carlo’ analysis (e.g. Uyeno, 1992; Clemson et al., 1995).

5. Processing of sensitivity analysis results

A great deal of information can be generated in sensitivity analysis, so much so that there is a risk of the volume of data obscuring the important issues (Eschenbach and McKeague, 1989). For this reason, the modeller must process and/or summarise the

information to allow decision makers to identify the key issues. The following sub-sections cover various possible methods for processing results of a sensitivity analysis, ranging from very simple to very complex. For many of the methods of analysis, I suggest possible layouts for graphs and tables. There are many other layouts which may be more suitable than these for particular purposes. A number of examples are drawn from my research in agricultural economics.

5.1. Summaries of activity levels or objective function values: one dimension

The simplest approach to analysis of SA results is to present summaries of activity levels or objective function values for different parameter values. It may be unnecessary to conduct any further analysis of the results.

A simple example of such a summary is presented in Fig. 1. This example (like several which follow) is from MIDAS, a linear programming model which selects optimal combinations of farming enterprises for representative farms in a region of Western Australia (Morrison et al., 1986; Kingwell and Pannell, 1987). Fig. 1 shows how the optimal area of wheat varies as a number of parameters are varied either side of their standard values. Each of the parameters in this example is varied up or down by amounts reflecting their realistic possible ranges. The format in Fig. 1 allows results from several parameters to be presented on a single graph. This allows easy com-

Fig. 1. Graphing changes in multiple parameters for a single output variable.

parison of the relative impacts of these parameters when varied over their realistic ranges, and these ranges are communicated by the horizontal span of the lines. In this example one can see that wheat yields have the biggest impact on the optimal area of wheat. Eschenbach and McKeague (1989) refer to this type of graph as a ‘spider diagram’, for obvious reasons. Another variation is to also plot the vertical axis in percentage terms so that the graph illustrates ‘elasticities’ (see Section 5.3).

Spider diagrams like these can also be constructed with the objective function value rather than an activity level as the dependent variable, allowing the decision maker to assess the sensitivity of the objective function value to parameter changes. For example if the objective is to maximise profit, this type of diagram reveals whether any parameter changes would result in a negative profit.

A potential problem with the use of percentage changes in spider diagrams is that if the parameter is small (e.g. variation is centred around zero), percentage changes may be large relative to those for other variables. In fact, if the initial parameter value is zero, percentage changes to the parameter are not defined. For these parameters, it may be appropriate to use an absolute change.

Spider diagrams are usually practical only for displaying the levels of a single activity. Where there are several important variables to display, one normally needs to limit results to changes in a single parameter. Fig. 2 is an example from MIDAS show-

Fig. 2. Graphing multiple output variables for changes in a single parameter.

Fig. 3. Graphing the allocation of a resource among alternative uses for changes in a single parameter.

ing production of wheat grain, lupin grain, pea grain, and wool as a function of wheat price. Because of the different scales of production, wool is shown on the right hand axis. This graph reveals that the main effect of increasing wheat price is to increase wheat production at the expense of wool. There are also smaller changes in the production of lupin grain and pea grain.

A different way of summarising the same model results is to show the allocation of a particular input or resource to the different possible outputs. The way these allocations vary can be effectively displayed by stacking the lines or bars, as shown in Fig. 3. This shows the allocation of land to production of each of the four products, with the allocations mirroring the trends in Fig. 2.

5.2. Summaries of activity levels or objective function values: higher dimensions

In Fig. 1, because all parameters but one were held constant for each line on the graph, it was possible to display results for several parameters on the same graph. In displaying the results of changing parameters simultaneously, it is difficult to handle more than two parameters in a graph without it becoming complex and difficult to follow. Fig. 4 shows an example of a method for displaying results from sensitivity analyses on two parameters. This

Fig. 4. Graphing combinations of parameter changes.

figure shows the impacts of changing wheat price and wool price on the optimal area of wheat selected by MIDAS. There are many other formats for three dimensional graphs which can be used for this purpose.

Results for more than two parameters require a series of graphs or a table. Well structured tables are probably the better option. Another approach is to develop an interactive database of model results, allowing decision makers to select the parameter values and displaying the corresponding optimal solution. This type of database acts as a simplified (and much quicker) version of the full model.

A final possible approach to the analysis of multi-dimensional sensitivity analysis is to use statistical regression techniques to fit a smooth surface to the results (Kleijnen, 1992, 1995b). This approach provides an equation which approximates the functional relationship between the parameter values and the dependent variable (e.g. the activity level or objective function value). Such an equation will be smoother than the step functions often produced by mathematical programming models and this may be useful for producing graphs or for conducting some of the analyses outlined below.

5.3. Slopes and elasticities

The rate of change (the slope) of an activity level or of the objective function with respect to changes in a parameter is an even briefer summary of the issue than the graphs shown so far. An issue is the need to compare slopes for different parameters. The

units of measure of different parameters are not necessarily comparable, so neither are absolute slopes with respect to changes in different parameters. One can often overcome this problem by calculating 'elasticities', which are measures of the percentage change in a dependent variable (e.g. an activity level) divided by the percentage change in an independent variable (e.g. a parameter).

$$\varepsilon = \% \Delta Y / \% \Delta X \quad (1)$$

or

$$\varepsilon = \delta Y / \delta X \cdot X / Y \quad (2)$$

A comparison of elasticities of an activity level with respect to different parameters provides a good indication of the parameters to which the activity is most sensitive. Table 2 is an example of such a comparison for MIDAS. The elasticities have been calculated assuming base values for parameters other than the one in question. Results have been smoothed using regression analysis and elasticities have been calculated from the fitted smooth curves.

An alternative to the use of elasticities is to standardise parameters as follows:

$$Z = (X - b) / a \quad (3)$$

where b is the base value for X and a is the range (i.e. $X_{\max} - X_{\min}$) (Kleijnen, 1995a).

5.4. Sensitivity indices

A sensitivity index is a number calculated by a defined procedure which gives information about the relative sensitivity of results to different parameters of the model. A simple example of a sensitivity index is the elasticity of a variable with respect to a parameter (Section 5.3). The higher the elasticity, the higher the sensitivity of results to changes in that

Table 2

Elasticities of optimal wheat area with respect to changes in various parameters

Parameter	Elasticity of optimal wheat area
Wheat price	1.5
Wheat yield	1.4
Wool price	-0.5
Lupin price	-0.3
Machinery size	0.0

parameter. Hamby (1994) outlined 14 possible sensitivity indices for cases where only a single output variable is to be evaluated, including the ‘importance index’, the ‘relative deviation’ index, the ‘partial rank correlation coefficient’, the Smirnov test, the Cramer–von Mises test, and a number of others. These are not outlined in detail here because many of them are complex and time-consuming to calculate. Furthermore, Hamby (1995) conducted a detailed comparison of the performance of each of the indices relative to a composite index based on ten of them. None of the complex indices tested performed as well as a simple index proposed by Hoffman and Gardner (1983):

$$SI = (D_{\max} - D_{\min}) / D_{\max} \quad (4)$$

where SI is the sensitivity index, $D_{\max}$ is the output result when the parameter in question is set at its maximum value and $D_{\min}$ is the result for the minimum parameter value.³ In cases where comparisons between different models are not important, the following even simpler sensitivity index can be perfectly adequate (and perhaps even preferable).

$$SI = (D_{\max} - D_{\min}) \quad (5)$$

Alexander (1989) suggested a number of complex indices for use in situations where the modeller wishes to assess the sensitivity of several output variables simultaneously. For example, for cases where the result of interest is a ranking of several variables, Alexander provides an index which indicates the sensitivity of the ranking to changes in a parameter.

5.5. Break-even values

Consider the question: “If parameter X were to change from its current value, by how much would it have to change in order for the optimal solution to change in a particular way?” This break-even approach addresses the issue of uncertainty about parameter values in a way which is often particularly

Table 3

Break-even changes in parameter values for cropping to be as profitable as pasture production on Soil type 1

Parameter	Break-even parameter change (%)
Wheat price	+ 50
Wheat yield on Soil type 1	+ 40
Wool price	– 80
Pasture yield on Soil type 1	– 70
Lupin price	+ 130
Lupin yield on Soil type 1	+ 120

helpful to decision makers. It helps in the assessment of whether the critical value of the variable falls within the range of values considered reasonable for the variable. If not, the decision maker can be advised (for the purposes of planning) to disregard the possibility of the variable taking a different value. If the break-even value is within the realistic range, this information can be used to justify collection of additional information to help predict the actual value of the parameter.

Table 3 shows an example from MIDAS. In the standard version of this model, the optimal use of land of a particular type (Soil type 1) is to grow pasture for grazing by sheep. The aim is to determine the circumstances in which cropping would be as good as or better than pasture. The table shows break-even percentage changes in various parameters (changes needed for the profitability of cropping on Soil type 1 to equal that for pasture). By judging whether parameter changes of at least this magnitude are ever likely to occur, the modeller can judge whether cropping is ever likely to be recommended on this soil type.

5.6. Comparing constrained and unconstrained solutions

The approaches discussed so far are based on assessing the sensitivity of the model to changes in parameters. A different approach is to add constraints to the model so that it is forced to adopt other interesting strategies. It is often very valuable to know how other strategies perform relative to the optimum. Fig. 5 shows an example, where the MIDAS model has been constrained to plant crops on various percentages of the farm area. Such a graph is

³ A referee suggested that where the relationship between X and Y is not monotonic, it may be preferable to use the highest and lowest value of D within the range for X , rather than the values of D corresponding to the extreme values of X .

valuable if the decision maker wishes to consider other strategies which achieve objectives other than that represented in the model. Fig. 5 shows how much profit must be sacrificed if the farmer wishes to deviate from the optimal cropping area of 60%.

A useful way of indicating the flexibility available to the decision maker is to report the set of strategies with objective function values within a certain distance of the optimum. For example, any area of crop between 40 and 70% of the farm is within \$5000 of the maximum profit. This is an example of one approach to testing the 'robustness' of a solution (one of the uses of SA listed in Table 1).

Sometimes it is useful to constrain the model to exclude an option in order to calculate the total contribution of this option to the objective, and to identify the best strategy which does not include it. Table 4 shows a summary of the MIDAS solutions which include and exclude the option of growing lupins on the farm. It is apparent that the inclusion of lupins increases profits by around 66%.

5.7. Using probabilities

A common characteristic of the methods of analysis presented above is that they do not require the modeller to explicitly specify probabilities of different situations. Sensitivity analysis can be extremely effective and useful even without taking this extra step to a more formal and complex analysis of results. In excluding probabilities from the analysis, the modeller is relying on the decision maker to give appropriate weight to each scenario. On the other

Fig. 5. Comparing optimal with sub-optimal strategies.

Table 4

Profit and optimal rotations with and without lupins

	Lupins included	Lupins excluded
Whole-farm profit (\$)	40870	24553
Rotation selected		
Soil type 1	PPPP	PPPP
Soil type 2	CL	PPPC
Soil type 3	CCL	CCCC
Soil type 4	CCL	CCCC
Soil type 5	CCF	CCF
Soil type 6	PPPP	PPPP
Soil type 7	CCF	CCF

C, cereal crop; P, pasture; L, lupins; F, field peas.

hand, an analysis using probabilities may be unnecessarily difficult and time consuming to conduct, and is likely to be more difficult to explain to the decision maker. The potential simplicity of sensitivity analysis is one of its attractions: an analysis which is not understood is unlikely to be believed. Depending on the importance of the issue and the attitudes and knowledge of the decision maker, the best approach to sensitivity analysis might not involve formal and explicit use of probabilities. Even if a probabilistic sensitivity analysis is to be conducted, a simpler preliminary analysis may be useful in planning the final analysis.

6. Overall strategies for sensitivity analysis

The techniques outlined above are a powerful set of tools for assisting a decision maker. However the modeller needs to avoid conducting sensitivity analysis in an aimless or mechanical fashion. The approach should be adjusted to suit the decision problem. As the analysis proceeds, the results obtained may lead to further model runs to test ideas or answer questions which arise. In a thorough sensitivity analysis, a number of the approaches suggested in the previous section might be used.

Within these broad guidelines, there are very many overall strategies for sensitivity analysis which might be adopted. Here are three systematic suggestions of overall strategies which are likely to be effective in cases where the analysis is used to help make a

decision or recommendation about the optimal strategy.

Strategy A (the most comprehensive) is as follows.

1. Select the parameters to be varied. Identify a range for each parameter which realistically reflects its possible range. For example, use maximum and minimum values, or an 80% confidence interval, but not a uniform 10 or 20% either side of the expected value. Also identify other possible discrete scenarios requiring changes to the model structure or formulation (e.g. changes in the objective to be optimised, inclusion of additional constraints).
2. Conduct sensitivity analyses for each parameter individually, using two parameter values (high and low or maximum and minimum). Conduct sensitivity analysis for each discrete scenario individually.
3. Identify parameters and discrete scenarios to which the key decision variables are relatively unresponsive, using one of the sensitivity indices presented in Section 5.4 (e.g. Eq. (5)).⁴
4. Exclude unresponsive parameters and scenarios from further analysis. For the remaining parameters, consider whether they are likely to have high positive, high negative or low correlation with each other. If it is intended to use probability distributions for random sampling of scenarios or for summarisation of results, estimate the distribution for each parameter and, for cases of high correlation, estimate the joint probability distribution. Possibly also estimate probabilities for the discrete scenarios selected in Step 1.
5. Design and conduct a modelling experiment which allows appropriately for combinations of parameter changes, paying particular attention to the

cases of high correlation between parameters. Possibly use Latin hypercube sampling (Clemson et al., 1995) or, if the number of combinations is manageable, a complete factorial design.⁵ Repeat this for each of the discrete scenarios individually, or if practical, for all combinations of the discrete scenarios.

6. Summarise results. For each key decision variable, calculate the values of a sensitivity index for all parameters and discrete scenarios, and rank them by absolute value.⁶ These results can be reported directly or used to select which parameters will be examined in graphs and tables (e.g. spider diagrams). This approach helps to prioritise the presentation of results which is essential to avoid an overload of graphs and tables. It also allows the decision maker to focus on important parameters and relationships. Calculate break-even parameter values for particular circumstances of interest.
7. On the basis of results so far, identify a tentative best-bet strategy and several others of interest. The other strategies might be chosen because they contribute to objectives other than those represented in the model, or because they are of personal interest to the decision maker. French (1992) suggested focusing on 'adjacent potentially optimal' alternative solutions, meaning strategies which are close to the base-case optimum and which would become optimal if parameters changed sufficiently. It is not necessary to limit the analysis to such a narrow set of strategies, although one should be mindful of the number of solutions required in the next step.
8. Repeat the experiment (Step 5) with the model constrained to each of the strategies (from Step

⁴ Steps 2 and 3 constitute a very simple version of what is commonly termed 'screening', meaning the identification of important factors for investigation in a later modelling experiment (e.g. Kleijnen, 1995a, 1996). Among the more sophisticated approaches to screening suggested in the literature are sequential bifurcation (Bettonvil and Kleijnen, 1996), iterated fractional factorial design (Andres, 1996) and replicated LHS design (McKay, 1995).

⁵ Alternatively an incomplete factorial design such as 2^{k-p} (Kleijnen, 1996) may be useful. If the model output is to be used for estimation of a statistical regression model, the cost of excluding runs in the 2^{k-p} design is likely to be small.

⁶ Commonly there are interactions between parameters in their impacts. To avoid reporting an unrepresentative value for a sensitivity index, it is desirable to calculate an expected value (or at least a simple average) for the sensitivity index over all scenarios modelled.

- 7). Summarise these results. Identify scenarios (if any) where each strategy is optimal. Calculate the cost of each strategy relative to the best-bet. Possibly repeat this with another strategy as the best-bet. At this stage the modeller may wish to use probability distributions to make probabilistic statements about the results.
9. Attempt to draw conclusions. It can be helpful to the analysts to focus their thinking by trying to couch conclusions in terms similar to one of the following examples.
 - (a) The optimal strategy is *X* in almost any plausible scenario, so *X* is a safe best-bet strategy.
 - (b) In some scenarios the optimal strategy is *X*, whereas in these other scenarios the optimal strategy is *Y*. If you can predict or identify the scenario, it is important to do the right strategy in the right scenario.
 - (c) In some scenarios the optimal strategy is *X*, whereas in these other scenarios the optimal strategy is *Y*. However, the cost of doing the wrong strategy is very low, so it is not very important to worry about doing the right strategy in the right scenario.
 - (d) In some scenarios the optimal strategy is *X*, whereas in these other scenarios the optimal strategy is *Y*. The cost of doing the wrong strategy when the decision maker should be doing *Y* is low, but the cost of doing the wrong strategy when the decision maker should be doing *X* is high, so if it is not possible to predict or identify the scenario, *X* is a safe best-bet strategy.

These conclusions correspond to the following recommendations: (a) do *X*, (b) do either *X* or *Y* depending on specific circumstances, (c) do either *X* or *Y*, it does not matter which, (d) if in doubt, do *X*. In addition there is a converse set of conclusions about which strategies are never likely to be optimal: (e) never do *Z*, (f) in certain circumstances do not do *Z*, (g) if in doubt, do not do *Z*. Try to identify which of the categories (a) to (d) the problem falls into and whether it is possible to specify any strategies like *Z* in categories (e) to (g).

Strategy B (slightly less comprehensive) includes all of the steps of Strategy A except 7 and 8.

Strategy C (the simplest strategy which is still systematic and useful) includes only Steps 1, 2, 3, 6 and 9.

7. Reporting results of sensitivity analysis

It is common for written reports of sensitivity analyses in published papers to address only a subset of the issues on which the SA can provide information. Of course one must be selective in the reporting and discussion of results, but too often discussions of sensitivity analyses drift away from the central issue being investigated onto interesting but relatively unimportant details. In other cases, SA results are presented without sufficient discussion of their consequences and implications for the central issue. To avoid these traps, the following report structure is recommended as a standard minimum.

1. From the base-case model, or other information, what is the initial optimal recommendation which is to form the standard for comparisons in the SA?
2. Which parameters most affect the optimal recommendation? A table of values for a sensitivity index ranked according to their absolute value is recommended. If appropriate, what are the break-even levels of parameters for changes in the recommendation?
3. How does the optimal recommendation change if the important parameters (from 2) change?
4. What are the consequences of not following the optimal recommendation? For example, how much less profitable are other recommendations?
5. Overall, what level of confidence can there be that the recommendation is in fact optimal?

In addressing these issues, the space devoted to each need not necessarily be large, and the relative importance of each will depend on the particular study. Point 1 is particularly important as it ensures that the discussion of the SA will be well focused and relevant. The recommendation to state the 'level of confidence' is not intended to provoke a formal probabilistic or statistical statement, but at least some relatively informal and subjective statement of confidence should be made. If the conclusion is subjective, say so.

Avoid the trap of overloading the report with the results of Category 3. As noted above, a helpful strategy in this regard is to demonstrate that certain parameters have little impact on the important decision variables, and then to avoid reporting further results for these parameters.

8. Concluding comments

There is clearly much more to the use of a normative model than finding a single optimal solution. That solution should be viewed as the starting point for a wide ranging set of sensitivity analyses to improve the decision maker's knowledge and understanding of the system's behaviour.

Even without undertaking the relatively complex procedures which explicitly involve probabilities in the sampling of scenarios or interpretation of results, sensitivity analysis is a powerful and illuminating methodology. The simple approach to sensitivity analysis is easy to do, easy to understand, easy to communicate, and applicable with any model. As a decision aid it is often adequate despite its imperfections. Given its ease and transparency, the simple approach to SA may even be the absolute best method for the purpose of practical decision making.

Acknowledgements

The paper has been improved by detailed reviewer comments, for which I am grateful. I also thank Jack Kleijnen for making me aware of a set of literature I had not discovered. Much of the work for this paper was done while I was on study leave at the Department of Agricultural Economics, University of Saskatchewan, Canada and I thank the members of that Department for their hospitality. Members of seminar audiences at the University of Western Australia and the University of Melbourne also made helpful suggestions.

References

- Alexander, E.R., 1989. Sensitivity analysis in complex decision models. *J. Am. Plann. Assoc.*, 55: 323–333.
- Andres, T.H., 1996. Sampling methods and sensitivity analysis for large parameter sets. *J. Stat. Comput. Simulation*, in press.
- Baird, B.F., 1989. *Managerial Decisions Under Uncertainty*, An Introduction to the Analysis of Decision Making. Wiley, New York.
- Bettonvil, B. and Kleijnen, J.P.C., 1996. Searching for important factors in simulation models with many factors: sequential bifurcation. *Eur. J. Oper. Res.*, 96: 180–194.
- Canova, F., 1995. Sensitivity analysis and model evaluation in simulated dynamic general equilibrium economies. *Int. Econ. Rev.*, 36: 477–501.
- Clemson, B., Tang, Y., Pyne, J. and Unal, R., 1995. Efficient methods for sensitivity analysis. *Syst. Dyn. Rev.*, 11: 31–49.
- Dungan, D.P. and Wilson, T.A., 1991. Macroeconomic effects and sensitivity analysis. *J. Policy Modelling*, 13: 435–457.
- Eschenbach, T.G. and Gimpel, R.J., 1990. Stochastic sensitivity analysis. *Eng. Econ.*, 35: 305–321.
- Eschenbach, T.G. and McKeague, L.S., 1989. Exposition on using graphs for sensitivity analysis. *Eng. Econ.*, 34: 315–333.
- Fiacco, A.V., 1983. *Introduction to Sensitivity and Stability Analysis in Nonlinear Programming*. Academic Press, New York.
- French, S., 1992. Mathematical programming approaches to sensitivity calculations in decision analysis. *J. Oper. Res. Soc.*, 43: 813–819.
- Hall, N. and Menz, K., 1985. Product supply elasticities for the Australian broadacre industries, estimated with a programming model. *Rev. Mark. Agric. Econ.*, 53: 6–13.
- Hamby, D.M., 1994. A review of techniques for parameter sensitivity analysis of environmental models. *Environ. Monit. Assess.*, 32: 135–154.
- Hamby, D.M., 1995. A comparison of sensitivity analysis techniques. *Health Phys.*, 68: 195–204.
- Harrison, G.W. and Vinod, H.D., 1992. The sensitivity analysis of applied general equilibrium models: completely randomised factorial sampling designs. *Rev. Econ. Stat.*, 74: 357–362.
- Hoffman, F.O. and Gardner, R.H., 1983. Evaluation of uncertainties in environmental radiological assessment models. In: J.E. Till and H.R. Meyer (Editors), *Radiological Assessments: A Textbook on Environmental Dose Assessment*. US Nuclear Regulatory Commission, Washington, DC, Report no. NUREG/CR-3332, pp. 11.1–11.55.
- Kingwell, R.S. and Pannell, D.J. (Editors), 1987. *MIDAS, A Bioeconomic Model of a Dryland Farm System*. Pudoc, Wageningen.
- Kleijnen, J.P.C., 1992. Sensitivity analysis of simulation experiments: regression analysis and statistical design. *Math. Comput. Simulation*, 34: 297–315.
- Kleijnen, J.P.C., 1995a. Sensitivity analysis and optimization: design of experiments and case studies. In: C. Alexopoulos, K. Kang, W.R. Lilegdon and D. Goldsman (Editors), *Proceedings of the 1995 Winter Simulation Conference*. pp. 133–140.
- Kleijnen, J.P.C., 1995b. Sensitivity analysis and optimization of system dynamics models: regression analysis and statistical design of experiments. *Syst. Dyn. Rev.*, 11: 1–14.
- Kleijnen, J.P.C., 1996. Five-stage procedure for the evaluation of simulation models through statistical techniques. In: J.M. Charnes, D.J. Morrice, D.T. Brunner and J.J. Swain (Editors), *Proceedings of the 1996 Winter Simulation Conference*. In press.
- Lomas, K.J. and Eppel, H., 1992. Sensitivity analysis techniques for building thermal simulation programs. *Energy Build.*, 19: 21–44.
- McKay, M.D., 1995. Evaluating prediction uncertainty. Los Alamos National Laboratory, NUREG/CR-6311 (LA-12915-MS).
- Morrison, D.A., Kingwell, R.S., Pannell, D.J. and Ewing, M.A., 1986. A mathematical programming model of a crop–lives–stock farm system. *Agric. Syst.*, 20: 243–268.

- Nordblom, T., Pannell, D.J., Christiansen, S., Nersoyan, N. and Bahhady, F., 1994. From weed to wealth? Prospects for medic pastures in Mediterranean farming systems of northwest Syria. *Agric. Econ.*, 11: 29–42.
- Rios Insua, D., 1990. Sensitivity Analysis in Multi-Objective Decision Making. *Lecture Notes in Economics and Mathematical Systems* No. 347, Springer Verlag, Berlin.
- Sobieszczanski-Sobieski, J., 1990. Sensitivity analysis and multidisciplinary optimization for aircraft design: recent advances and results. *J. Aircr.*, 27: 993.
- Tzafestas, S.G., Theodorou, N. and Kanellakis, A., 1988. Recent advances in the stability analysis of multidimensional systems. *Inf. Decis. Technol.*, 14: 195–211.
- Uyeno, D., 1992. Monte Carlo simulation on microcomputers. *Simulation*, 58: 418–423.