

Sensitivity Analysis Revisted

By

J.E. Raiswell

Department of Management

Teesside Polytechnic

Abstract

The different approaches to measuring parameter sensitivity in system dynamics models are discussed and a practical evaluation attempted. For certain types of model behaviour, a procedure is proposed for generating performance measures which converge to values independent of simulated time. Using these types of measure in conjunction with local sensitivity analysis, most likely values of the performance measures may be estimated and then used to calculate the most likely value of performance improvement following system redesign.

1. Introduction

Most system dynamics modelling work has two distinct aims:

- (i) to understand the observed behaviour of the existing system
- (ii) to redesign the system for improved behaviour

(Coyle 1977)

In this context the term 'improved behaviour' frequently involves the attempt to make performance of the redesigned system robust against the likely external shocks to which the system may be exposed (external robustness) and also against the effects of any errors in the modelling process itself (internal robustness). Errors of this type may arise through uncertainties in initial conditions, parameter estimates or system structure. Sharp (Sharp 1974) has shown that each of these three types of error may be represented as errors in parameters of the model, and henceforth we shall use the term 'parameter errors' to refer to this group of errors.

In understanding behaviour of the existing system, our concern is to locate those feedback loops within the model to which system behaviour is most sensitive. Despite the complexity of many models of real systems it is frequently found that only a few of the feedback loops actually affect behaviour materially at any given time. Knowledge of these loops provides us with a basis both for understanding and for altering behaviour. In his search for these important areas of the model the modeller may use the techniques of either loop analysis or sensitivity analysis.

With its focus on the individual feedback loop, loop analysis is potentially the more useful of the two approaches. However with its present state of development, the technique suffers from the important disadvantage that it does not yet permit exhaustive analysis of complex models. Sensitivity analysis however does allow exhaustive analysis of the effects of parameter changes on each of the performance measures of the model. Whilst the two techniques would probably be used together in a study of any particular model, sensitivity analysis seems to be the more simple to apply and also seems to offer further advantage through its exhaustive nature.

For the second aim of system dynamics modelling, we may regard the concept of internal robustness as a lack of sensitivity of system performance to the types of modelling uncertainty we have mentioned. Clearly therefore a second sensitivity analysis is required to check the performance of the redesigned system.

2. Types of Sensitivity Analysis

We now distinguish two broad classes of sensitivity analysis each of these in turn splitting into two subclasses. Focussing attention exclusively on parameter sensitivity the major split conceptually is between dynamic and point sensitivity analysis.

The value of a given performance measure P_i is a function of time t and each of the model parameters S_1 to S_n , $P_i = P_i(t, S_1, \dots, S_n)$. In dynamic sensitivity analysis the sensitivity coefficients ($\partial P_i / \partial S_j$) are calculated for each time step throughout the simulated time. From the point of view of system understanding, such an analysis allows us to detect the changing importance of different parameters (and hence different loops) throughout the period of simulation. A good example of this type of analysis is shown by Burns and Malone (1974). Although the computer time required to perform this type of analysis can be quite large, the major practical disadvantage of dynamic analysis is the very large volume of information generated. As we have discussed, the modeller's main interest is in knowing which loops dominate behaviour at different times in the simulation. Consequently most of the information produced is of rather little value.

In point sensitivity analysis, the sensitivity coefficients ($\partial P_i / \partial S_j$) are calculated only at a single point in time, usually the end of the run. Since the coefficients are in fact functions of time, point sensitivity analysis gives only a snapshot of the parameters affecting model behaviour. The volume of information produced is however quite manageable.

Within either class of analysis we may consider either local sensitivity or global sensitivity. The distinction between these lies in the magnitude of the parameter changes to be investigated. Within local sensitivity analysis parameter changes are constrained to be small, typically about 1% of the parameter value itself. With this restriction, the effects of small changes made one at a time to each parameter are measured. The effects of simultaneous changes to parameters may then be estimated by adding together the individual effects. Local sensitivity analysis of this type provides a natural starting point for any analysis of model behaviour, principally because it is straight forward to apply and to interpret the results.

Global sensitivity analysis is concerned to measure the effects of large changes to parameters. Because of synergistic effects within the system, the effect on the model of global changes made to several parameters simultaneously may not be estimated by summing their individual effects. Global sensitivity is used principally to test system behaviour under extreme conditions and as a final test of robustness.

In summary then, current techniques permit either point or dynamic sensitivity analysis with either local or global changes made to parameters. Of these point, local analysis is the most straight forward to apply but, unless care is taken in the design of performance measures, is least useful. A similar criticism may also be applied to point, global analysis. Dynamic analysis on the other hand provides such large volumes of information, that the modeller may have difficulty in selecting the information he needs. Two lines of development are therefore suggested:

- (i) to attempt to overcome the basic limitation of point sensitivity analysis that it relates only to a single point in time
- (ii) to develop post-processors for dynamic sensitivity analysis, so that only the most important information is presented to the modeller

In this paper we shall focus on the first problem.

3. Design of Performance Measures for Point Sensitivity Analysis

As we have discussed point sensitivity analysis provides a snapshot of the coefficients $(\partial P / \partial S_j)$ for all i and j at a given time. If the performance measures used in the model are based on instantaneous or short time averages of model variables, the coefficients calculated change with time and are little value in understanding model behaviour at times other than when the snapshot was taken.

An alternative to this is to define performance measures which tend to converge towards fixed values independent of time. If this can be done study of these sensitivity coefficients will enable the modeller to gain knowledge of the important parameters over the whole simulation. One way in which such performance measures may be constructed is to take the arithmetic mean of key variables over the whole simulation. If these key variables, show repetitive, or constant behaviour the mean value will simply tend towards a constant, whose value is independent of time.

More precisely if X is a model variable which after an initial period settles down to an oscillatory behaviour pattern about a constant value (or is itself constant over this latter part of the simulation) then provided T is large relative to the initial period then the measure $P = \frac{1}{T} \int_0^T X(t) dt$ will tend towards a fixed value as T is increased. For this type of system a point sensitivity analysis in which $(\partial P / \partial S_j)$ is evaluated at time T can give useful information about the sensitivity of P to parameter changes over the whole simulation.

The approach we are advocating, is clearly restricted in its applicability but we would argue that the class of models for which it is useful is nevertheless quite large. If such performance measures may be constructed in a given situation, and are indeed useful knowledge of the sensitivity coefficients found may be put to further use. Before developing these ideas however we turn our attention to some practical aspects of the sensitivity analysis of more general applicability.

4. Conducting the Sensitivity Analysis

In any parameter sensitivity analysis we are interested in locating the most important parameters to which the performance measures are most sensitive. The situation is likely to be one in which we have say 10 or 20 performance measures and possibly hundreds of parameters, the number of sensitivity coefficients calculated may therefore run into thousands. Since the values of the performance measures and the parameters may differ by many orders of magnitude it is not obvious from study of the sensitivity coefficients, which the most important parameters are.

To overcome this problem, it has been found useful to calculate the relative elasticity of each performance measure with respect to each parameter change, in addition to the sensitivity coefficients.

The relative elasticity of a performance measure P_i with respect to a parameter S_j is RE_{ij} where this is defined as $RE_{ij} = (\partial P_i / \partial S_j) (S_j / P_i)$.

The expression is clearly dimensionless and its value may be used quickly to spot those parameters affecting the performance measure. Values of RE_{ij} lying between +1 and -1 indicate those parameters for which a given fractional change in S_j gives rise to a lower fractional change in P_i . Generally these are the parameters which may be ignored. Values outside this range may be used to direct attention to those parameters significantly affecting the performance measure.

5. Use of Local Point Sensitivity Analysis: Most Likely Values of Performance

Returning now to the more special case of local point sensitivity analysis, let us suppose that a performance measure has been defined as in section 3 and that a point local sensitivity analysis has been carried out so that the set $(\partial P / \partial S_j)$ is known for all j . Now assume that the distribution of likely errors in each of the S_j is known and that these errors are small in relation to the S_j . Then we can regard each parameter value used in the base model, upon which the measurements were made, as being a single sample from the known distributions. Taking different samples from each of these distributions in conjunction with the sensitivity coefficients, we may estimate the effect on the performance measure, simply by adding the individual effects of these parameter changes. By repeating this process many times in a monte carlo simulation, we may build up a distribution of possible values of the performance measure.

In essence this process involves repeated use of the relation:

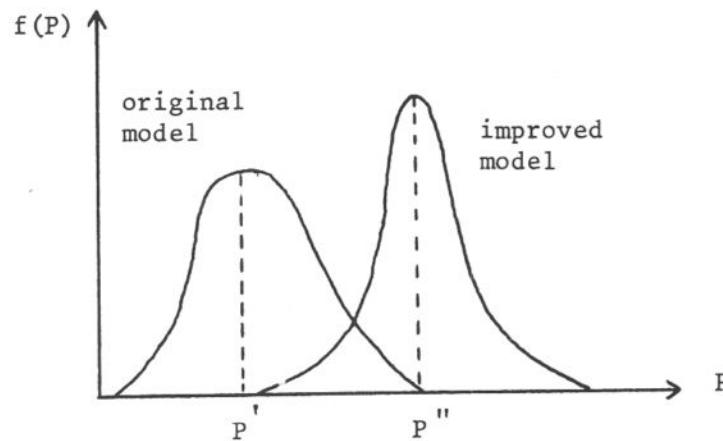
$$P \cong \sum_j \left(\partial P / \partial S_j \right) \Delta S_j$$

where each of the ΔS_j used, is taken from the assumed distribution of errors in the S_j .

Generally it will be found that the value of P generated by the base model is not the model value of the distribution found by this process. Since the base model may be regarded as a single sample from the distribution of P , we may interpret the model value of P as the most likely value of P local to its original value.

If now this whole process is repeated after system redesign, we can construct a new distribution for the same performance measure.

Putting these two distributions together we obtain:



P' and P'' are the most likely values of P resulting from the original and improved models respectively; an estimate of the most likely value of the improvement made by system redesign is therefore $(P'' - P')$.

Now if each of the distributions of P is normalised so that the area under each curve is unity, then $f(P)$ becomes the probability density function for P in each of the two cases. Hence if we regard real world parameter values as giving rise to random samples from each of the two distributions we may use the curves to estimate the probability that a performance improvement is actually obtained.

Use of these ideas therefore permits the modeller to make rather more precise statements about the likely magnitude of improvements through system redesign and also the chance that these will be achieved in practice.

6. Conclusions

It is hoped that the ideas put forward in this paper will encourage modellers to make more systematic and extensive use of the concepts of local sensitivity analysis. In a given situation, if performance measures of the type described can be constructed, sensitivity of these measures to parameter changes can lead directly to a greater understanding of system behaviour. As we have shown, use of these ideas in conjunction with estimates of parameter errors can be used to estimate most likely values of performance measures and hence to estimate the chance that the predicted improvements will be realised in practice. Where applicable, the approach described can be used to avoid some of the charges of arbitrariness in selection of parameter values, frequently levelled against system dynamics workers.

References

1. Burns, J.R. & Malone, D.W. (1974) 'Computational Techniques for Analysis of System Dynamics Models of Social Systems' Simulation Planning Science, Vol.X, pp215-223.
2. Coyle, R.G. (1977) "Management System Dynamics" John Wiley & Sons
3. Sharp, J.A. (1974) 'Some Problems of System Dynamics Methodology' Unpublished Ph.D. thesis, University of Bradford.