

Practical Industrial Application of
System Dynamics

R.G. Coyle
System Dynamics Research Group
University of Bradford
U.K.

Abstract

Since the publication of Industrial Dynamics there have been very few published accounts of the industrial application of the methodology. Indeed, the main use of the technique has been in large models of socio-economic problems. This paper discusses the practical application problem, and describes several cases of real industrial studies. It puts forward general criteria for successful practical studies, which apply to socio-economic as well as industrial problems.

Introduction

The original source of system dynamics, Forrester (1961), devotes considerable space to a model of an industrial company. In many places in the book Forrester persuasively argues that system dynamics methods are applicable to industrial problems. Despite this, there have been very few real industrial applications of SD modelling and the technique is taught at very few of our Universities.

The reasons for this state of affairs have been examined elsewhere (Coyle 1971 and 1973) and may briefly be summarised by saying that it is very hard to see from 'Industrial Dynamics' how one applies the method in practice. Therefore very little use is made of it and so few management science practitioners have any idea of how useful it can be and instead tend to think in terms of its obvious and real limitations. This is an unfortunate situation because the main ideas of SD are very appealing, particularly in relation to two types of managerial problem.

The first of these is the control of a system in the face of shocks from a non-stationary stochastic process. The classical example is a production/Distribution system for consumer goods sold from stock. In such a situation 'demand' (and the terms have to be defined precisely in each problem) exhibits day-to-day variation which may well be random, but there is also the possibility of genuine movements in the general level of demand. Clearly one wishes to make the system insensitive to noise but responsive to general trends. This is a standard problem in control theory and SD is well suited to its solution.

It is worth noting that conventional stock control theory gives very poor solutions to this problem because it introduces high gain into the system which usually has the effect of making the dynamics of the production side of the system far more severe than they need to be. The correct approach is to alter the bandwidth of the system to reduce its noise sensitivity.

The other general area is in corporate planning, the very object of which is to introduce dynamic behaviour into the system. In such a case one would naturally appeal to dynamic analysis in the search for conflicting corporate policies and the amelioration of their effects.

Both of these situations usually involve forecasting by the system and it is popularly believed that more 'accurate' forecasting will improve the behaviour of the system. The interesting point is, however, not the forecasting itself but also what the system does with the forecast when it is made. This raises important, and interconnected, questions such as which variables should be forecast, how they should be responded to, how much uncertainty in forecasts the system can tolerate and how the time constants of the system are related to the period for which the forecast should be made. Again, dynamic analysis (as opposed to DYNAMO simulation) is the proper approach to this problem area.

Problems of Practical Application

The research programme of the Bradford group has been aimed at expanding the practical application of SD, though we imply no claim to be the fount of all knowledge and recognise the debt we owe to many other workers. Broadly, however, we have attempted to work on three problems which have seemed to us to inhibit widespread practical use of SD.

- a) It is hard to get started on a dynamic model, and even harder to stop.
- b) The methodology is in some ways, weak.
- c) There is a great shortage of examples of what other people have done to guide the aspiring user.

The third area is dealt with more fully in the body of this paper. The others are extensive areas in their own right, but must be dealt with cursorily here.

In order to aid the starting and stopping problem two guideline methods (it would be inappropriate to call them techniques) have been evolved. These are list extension and type assignment (Coyle 1974). These methods depend on the use of influence diagrams, rather than the Forrester-type flow diagram. List extension is an approach which starts with the analyst nominating a few variables which are the object of this study, and guides the modelling process by stopping whenever a closed set of feedback loops has been found. If this model is inadequate the process continues until the next larger set of loops emerges and so on until the model is deemed to be satisfactory. Type assignment determines the numbers of levels and rates which the model should have and provides important checks on the validity of the formulation.

Methodologically, too little is known of the mathematics of SD models and the DYNAMO language, while convenient, lacks many useful facilities and allows some rather unjustifiable model formulations. Sharp (1974), has tackled the mathematical complexities and has produced some useful results on sensitivity and aggregation.

We have also developed an improved DYNAMO-type simulation language, known as DYSMAP (Ratnatunga 1974). This provides a variety of facilities not easily available in DYNAMO. In particular it permits the inclusion of FORTRAN macros in the DYNAMO program and produces a FORTRAN translation of DYNAMO statements which can then be operated and adapted as a FORTRAN program using pre-written routines for output.

Industrial Case Studies

The main aim of this paper is to describe actual applications of SD in British industry. Only a summary can be presented here and only some of the cases can be mentioned. In passing, it is worth pointing out that most of these projects did not involve DYNAMO-type programming at all.

Hughes (1971) worked for Cadbury's on the problem of production and sales forecasting for their Christmas product line. Production of Christmas products must commence in May and, at any time, can be curtailed or stopped there being, of course, effects on the amount which can then be produced during the remainder of the year. Orders are taken by the salesmen, starting in April or May and building up to a large volume in October/November. The company also has the option of forbidding salesmen to take any further orders, but this is a drastic measure which is rarely invoked. The problem, therefore involves deciding how to respond to uncertain forecasts soon enough to be able to do something about it. It can be formulated for dynamic programming, but the result is intractable and, in some ways misses the point which is that one has a large degree of freedom to change the control structure in business problems, and this is usually a very profitable avenue to follow. Hill (1973) and Burdon (1974) have studied the dynamics of the Chemical Plant Investment Cycle and this is to be expanded by employing an SSRC Research Fellow for two years. The work has attempted to explain and model the cyclical behaviour of the industry and to suggest ways of reducing it. The interesting aspect, and one of the factors which makes SD such a challenging area, is that there are really three problems. The first is the general, industry wide problem. The second and third are how should a chemical company, or a plant manufacturer, best live with the cycles and can they, even, profit from them. Whilst being anti-social both of these are more realistic problems than the first. Clearly, they arise in practically any industry.

Coyle (1974, Chapter 10) and Sharp and Coyle (1974), have worked on problems of production and distribution in a national consumer-goods producer which manufactures in the U.K. for the home and EEC markets. The first study treated the way in which demand and inventory policy for EEC sales was translated into demand on the UK manufacturing facilities. It was shown that the way in which production was ordered introduced very large dynamics into the production system; far larger, in fact, than those in the market. An alternative ordering policy smoothed the load on production, and could tolerate much larger variation in the market demand with lower planned stock levels.

The second study was directed at the detailed control of production for both UK and EEC markets. The most interesting feature was order backlog and there was considerable argument in the company about the extent to which larger backlogs represented genuine business as opposed to customer ordering to get a place in the queue. An analysis of the system showed that alternative production strategies reduced the backlog so that it was irrelevant what one believed about whether it was real. It was also shown that the real problem was not production control but capacity control.

Barnett (1973) analysed policies for the development of oil fields in the light of uncertain information about the size of the field and the productivity of wells drilled into it. The main problem is that the petroleum engineer must estimate the field parameters on the basis of his judgment of the similarity between the geology of the new field and others which he knows of. In general he tends to be pessimistic in his estimates and Barnett showed that this was the proper attitude.

This has an interesting connection with the forecasting problem mentioned earlier. Barnett also showed that forecasts of oil price had no bearing on the exploitation of a field. This is a very surprising result which undoubtedly calls for confirmation but is, again, in line with our experience of forecasting and system dynamics, namely that forecasting of 'obvious' variables may degrade system performance by creating variable gains in important feedback loops.

Two other studies may be mentioned as they bear on the wider applications of corporate systems.

An investigation has been conducted into the supply system of a major oil company. This system embraces all the activities of oil production, its transport to the market area, the holding of crude stocks and refining and product stock holding. The study cannot as yet be reported in detail but some rather surprising results have emerged. Probably the greatest benefit in this case was a learning vehicle for the company, both in terms of knowledge of system dynamics as a technique and in understanding of the way in which the company is put together.

It is interesting to note the potential application of SD as a design tool in this case. The oil crisis of October 1973 has prompted a review of oil company operating procedures and the design creation of new control mechanisms. The potential of SD for the rapid design of such systems, rather than simply allowing them to evolve in the course of time is considerable. (Coyle, 1974, Chapter 11).

Finally we refer to a project which is still in its early stages. A metals company refines ore into ingots, carries out initial processing known as semi-fabrication, and finally manufactures complete end-use goods for sale. The firm buys some of its ingot to overcome a shortage of refining capacity and sells a good deal of its semi-fab output to its own competitors because of its small share of the end-use market. Such a firm can grow by some combination of upstream integration into smelting, lateral integration of semi-fabrication, or downstream integration into manufacturing. The problem is to find a set of capacity-acquisition and financial control policies which will enable the firm to generate the finance needed to sustain its expansion, in the face of uncertainties about the future size of the market. Preliminary experiments with a very simple model indicate that substantial differences in performance are generated by different attitudes to expansion and liquidity. This is undoubtedly

a very promising area for future SD work.

Perhaps one common thread connecting these different projects is the small amount of effort needed to produce a significant result and the high degree of management acceptance of the ideas involved. The other unifying factor is that none of these studies, except Barnett's had anything to do with 'optimisation'. Not surprisingly we find that SD models contain so many indications of performance that optimisation would be so arbitrary as to be pointless.

Criteria for Successful Application

It would be false to imply that all these studies have been uniformly successful, or, if successful technically have necessarily been acceptable to management. At the risk of stating the obvious we can, however, draw on the experience to suggest some guidelines for successful implementation of SD.

a) Problem Identification

The client should be aware of a problem in which there is dynamic behaviour - either in the inputs to the system (e.g. demand) or in the state of the system (e.g. production level, unstable, cash flow variable, capacity acquisition).

The dynamics have to be identifiable, even though, in general, it will not be possible to find accurate statistical evidence for them.

The dynamics have to matter in that their existence creates a problem, or the firm is trying to bring them about (e.g. corporate growth).

There has to be at least the conceptual possibility of change to bring about improved control. Obviously if everything is inviolate there is no scope for improvement.

b) Feasibility

In nearly all cases it is a good idea to do a feasibility study to see whether the problem can be defined and solved, in principle, and, hopefully, in practice. This should aim at the following questions.

What is the problem and why does it matter? What, therefore, is the purpose of the study?

The question of the object of the exercise is of crucial importance and we believe that the following guidelines are useful:-

- i) Avoid the idea that anything will be 'optimised'. It is practically meaningless in a dynamic situation.
- ii) Aim instead at improving the behaviour modes of the system, e.g. making its response to a step change in demand more damped, or ensuring that a system in a growth mode does not go into overshoot and collapse.

- iv) Aim to make the system robust in the sense that it performs pretty well for every conceivable input, rather than superbly for the most likely situation.

Having defined the problem, one can then turn to more strictly technical aspects.

Can we build a simple influence diagram showing the presence of at least one feedback loop? What are the likely difficulties of programming and data collection?

What are the possibilities for alternative control policies, as indicated by the properties of the simple initial model? What are the behavioural problems of implementation?

c) Project Management

We have found that it is fairly easy to enlist management collaboration, partly because the influence diagram is something they understand, but also because SD seems able to tackle the problems of growth and overall management which they feel to be important. However, this attitude requires careful fostering and we do this by means of regular interim reports. The aim of these is to say, in effect, 'this is what we have found out so far, is it right?' This usually brings out a good deal of comment and the reports gradually become firmer and the models more satisfactory. Since the reports are about models we have found it valuable to include model output in the reports, even from the very simplest initial models. This is in line with the idea that a model should evolve, and should guide its own development.

Conclusion

To summarise, we believe that we have demonstrated that SD is a methodology which can be applied to real practical problems and that it forms a useful addition to the management scientists armoury. A good deal remains to be done, however, mainly along the lines of improving the techniques, better training facilities and teaching materials, and, finally, more published case studies.

References

- | | |
|----------------|------------------------------------------------------------------------------------------------------------------------------|
| Barnett, A.B., | A System Dynamics Model of an Oilfield's Development. Ph.D thesis, University of Bradford, 1973. |
| Burdon, A.J., | System Dynamics applied to a stock control problem in the Textile Industry. M.Sc dissertation, University of Bradford, 1974. |
| Coyle, R.G., | Management System Dynamics. Forthcoming from John Wiley. Preliminary edition by University of Bradford, 1974. |
| Coyle, R.G., | On the Scope and Purpose of Industrial Dynamics. International Journal of Systems Science, 1973, Vol. 4., No. 3. |

- Coyle, R.G., Some empirical evidence on the application of Industrial Dynamics. Unpublished paper of the University of Bradford, System Dynamics Research Group, 1971.
- Forrester, J.W., Industrial Dynamics, MIT Press, 1961.
- Hill, G.W., A dynamic model of investment by the Chemical Plant (and Petrochemical) Industry. M.Sc dissertation, University of Bradford, 1972.
- Hughes, K., Dynamic model of an industrial system to extreme seasonal cycles. M.Sc dissertation, University of Bradford, 1971.
- Ratnatunga, A.K., DYSMAP Users Manual. University of Bradford, 1974
- Sharp, J.A., A Study of some problems of Systems Dynamics Methodology. Ph.D thesis, University of Bradford, 1974.
- Sharp, J.A., and
Coyle, R.G., The Use of a System Dynamics Model for the Redesign of a Production Planning System. A case study. Unpublished paper, System Dynamics Research Group, University of Bradford, Revised 1974.