Preferred citation: R.H. Anderton. Industrial dynamics: a technique for gaining understanding of complex systems. In **Papermaking Systems and their Control**, *Trans. of the IVth Fund. Res. Symp. Oxford, 1969*, (F. Bolam, ed.), pp 633–647, FRC, Manchester, 2018. DOI: 10.15376/frc.1969.2.633.

INDUSTRIAL DYNAMICS: A TECHNIQUE FOR GAINING UNDERSTANDING OF COMPLEX SYSTEMS

R. H. ANDERTON, IPC (Group Management) Ltd., Feltham, Middx.

Synopsis Control of a complex organisation requires that its response to decision inputs be predictable. Traditionally, this has been made possible by almost breaking down the organisation into very simply coupled subsystems. This allows decision-makers to use a relatively simple and usually informal model.

Size and environmental rate of change are now tending to make the decoupling hard to make effective. Industrial dynamics provides both a theoretical structure and a computer simulation technique designed to provide decision-makers with a richer model of the dynamic feedback processes generally involved in industrial operations. The structure defines a system in terms of a closed boundary containing feedback loops, which in turn are formed by the relation between level and rate variables. The DYNAMO compiler allows models of systems conceived on these lines to be programmed in a straightforward way so as to allow experimental investigation of the system variables. The paper discusses case studies and concludes with notes on model validity and the problem of data.

Introduction

THIS paper was prepared as an introduction to the subject of industrial dynamics. It is thus very largely an account of ideas developed by Prof. Jay Forrester and his associates at M.I.T. to whose original writings the interested reader is referred.(1–2, 5, 8) We can start, however by establishing the need for such a technique in a particular context.

New kinds of help are needed for managers charged with the organisation and control of complex industrial systems.

The tendency is for these systems to become larger, more extensive and often more diverse. At the same time, they are being required to respond more rapidly and more flexibly to change, both in technology and in the pattern of demand for their output. Although these tendencies increase the amount of

Under the chairmanship of M. I. MacLaurin variety with which the manager must cope, there are rather definite limits to the amount of variety that an unaided human being can process. These limits are obvious enough in the straightforward case of sequential information processing such as reading, although the limits are rather high. They may be less obvious and more serious when this serial approach breaks down and relations between pieces of information are the essence of the matter. This occurs in risky situations when probabilities of separate events have to be compounded to estimate the probability of occurrence of a complex of events. It also occurs when the dynamics of situations determine the system behaviour, that is, when there are many strong couplings between parts of the system showing relationships developing over time.

In both these situations, the predictive power of the unaided human brain is rather weak. The matter is made worse because typically the weakness is not understood. It is the argument of this paper that—

1. These pressures on our mental powers are becoming increasingly felt both directly in our industrial systems and indirectly through failures in control in the economic and social systems that impinge on them.
2. The techniques of industrial dynamics developed by Forrester and his co-workers can provide powerful help in tackling the problems created by these pressures.

A managerial dilemma

IT WILL be interesting first to look at some of the ways in which the impact of these problems have been reduced by organisational conventions. We have said that there is a rather low limit to the number of interelationships we can handle together. (It seems that we have difficulty in predicting the behaviour of systems in which more than four or five variables dynamically interact.) As we have to have at least some competence in predicting system behaviour, if we are, in any sense, to control it and as industrial systems contain typically many more than five significant variables, we have little choice in organisational design. The subsystems must be largely decoupled so that the interactions between parts are simple. Putting it another way, the system must be almost broken down. We can give three examples of how this is done—

1. The use of buffering stocks to separate stages in a manufacturing process is a familiar one. This has a similar effect to that of a low pass filter in electrical engineering. Rapid changes are absorbed by the stock and only slow variations are transmitted. Thus, in the short term, the systems are decoupled.
2. The divisionalisation of large corporations by market or product or sometimes by function.

3. Less apparent, but of great importance, is the separation of activities according to the characteristic time that associated changes take to occur. Activities in an organisation can be thought of as forming a hierarchy in which higher levels have longer characteristic times.

   If these characteristic times are separated by factors of about 5–10, then the higher level is quasi-static from the viewpoint of the lower so that the higher level factors can be thought of as fixed policies or constraints.

These conventions have been (and are) remarkably effective in decoupling complicated systems and so making them manageable. They allow the manager to hold in his head a simple model that adequately represents the system he is managing.

Yet, we will now point out how, in each of the three examples given, current trends are making the solution less effective, in that the economic costs are increasing while the solution itself can at the same time be raising new dynamic problems.

TABLE 1—AN EXAMPLE OF THREE STOCKS IN CASCADE
(The figures are percentages)

	Retailers				*Wholesalers*				*Finished goods store*			
Week	*Out*	*In*	*Balance*	*On order from wholesalers*	*Out*	*In*	*Balance*	*On order from finished goods store*	*Out*	*In*	*Balance*	*On order from production*
1	100	100	100	100	100	100	100	100	100	100	100	100
2	110	100	100	130	100	100	100	100	100	100	100	100
3	110	130	110	110	130	100	70	190	100	100	100	100
4	110	110	110	110	110	190	150	70	190	100	10	370
5	110	110	110	110	110	70	110	110	70	370	310	Nil
6	110	110	110	110	110	110	110	110	110	Nil	200	20
7	110	110	110	110	110	110	110	110	110	20	110	110

Note; The order rule is to order sufficient to provide a balance equivalent to one week's sales as estimated from last week's sales

In the case of buffering stocks, the investment will increase if demand becomes more fluctuating, less predictable and more varied. If operations become more extensive, time lags and distortion in information and (possibly) naïve computer programming may tend towards the dynamic effect caricatured in Fig. 1.

In the case of divisionalisation, as products and markets change their identity, opportunities for synergy may be lost if rigid decoupling is maintained. Dynamically, too, inconsistent divisional targets can result in pendulum swings as, again, time lags and distortion in inadequate information links constrain the process of compromise.

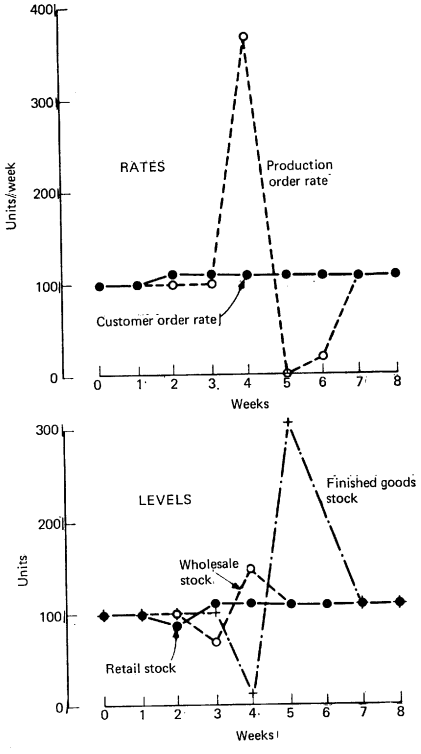

Fig. 1—Example of three stocks in cascade

In the third example, the time-separating hierarchy, the difficulties are more fundamental. The environmental rates of change are becoming such that the quasi-static assumptions of lower echelons are no longer valid.

For example, rapid external changes in market demand and in technology produce competitive pressures for internal changes at traditionally high levels of the hierarchy. These changes may be so frequent that the lower levels may not be able to increase the tempo of their activities in proportion and the time separation breaks down.

Without a stable policy framework within which to work, lower echelons are constantly attempting to adjust to or anticipate new situations. A steady state is almost never reached. A major decision input will now set off long chains of response and reaction, some links of which will feed back on themselves, positively or negatively. Intuitive predictions of results based on simple concepts of causality are now unlikely to be useful. Sometimes in fact the results will be counter-intuitive, nor is formalisation of the model through standard statistical and econometric methods likely to help. The presence of time lags, autocorrelation, non-linearity and large numbers of variables make these techniques extremely difficult to apply, to say the least of it.

What is needed is, firstly, a theoretical structure that provides a language and concepts, secondly, a simulation technique based on this theory that, together, will allow real complex systems to be modelled and explored. With these richer models, complexity can then be accepted and coped with, rather than artificially cut down.

A theory and technique that promises to meet these needs will be briefly described in the next section.

The industrial dynamics technique

The theoretical framework—Industrial dynamics provides a theory of system structure that can be outlined as follows. A system is a collection of parts

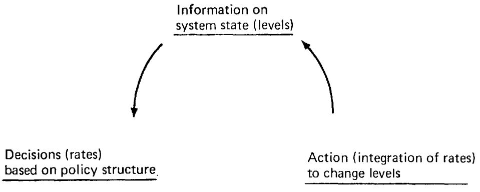

(a) Basic loop structure

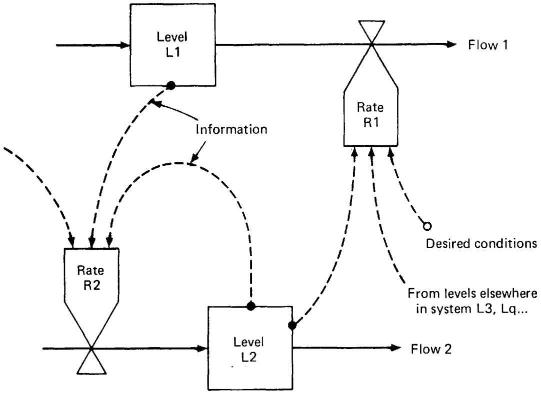

(b) Portion of a system in flow chart form

Fig. 2—System structure in industrial dynamics

connected to achieve a purpose. On this basis, a *system boundary* can be defined that includes the area of interest. System behaviour is generated inside this boundary. Within this boundary, the system is built up by the interaction of *feedback loops*. These are of two basic types, negative feedback loops (which are goal seeking) and positive feedback loops (which produce growth). These loops are constructed by connecting two types of variable. The first are called *levels* and are the accumulated quantities in a system (stocks, for example). The second are *rates* and fix the flow of quantities, which will be integrated to change a level. In a system, rates are determined by the 'policy' structure of the system, which in turn is defined by the goals and observed states, together with a procedure for converting the discrepancy between these two into a desired action. This system structure is illustrated in Fig. 2. This theory of system structure is held to be of very general validity, describing systems that extend beyond industrial ones(1, 2) to include social and economic,(3) even biological situations.(4)

What then becomes possible is a very direct modelling of hypotheses on the *mechanism* of the system under study. In particular, it stresses the role of information in integrating an organisation and enables the importance of its quality and timeliness to be evaluated. A clear distinction is made between the design of policy that determines how information is converted into action and particular decisions made within a given policy structure.

Simulation procedures—If a system is described in accordance with this scheme, it becomes a straightforward matter to write sets of difference equations that define the system operation through time. Numerical solution of these equations then allows the trajectory of the system to be plotted, starting from some initially defined states. Time is quantised into discrete solution intervals made small enough for interpolation between solutions to be accurate enough for the purpose in hand.

During each solution interval, *rates* are calculated using sets of *auxiliary* equations supplied with the values of the levels *existing* in the previous solution interval. New level values can then be arrived at by adding to the old ones the appropriate rate multiplied by the solution interval.

The DYNAMO compiler(5) enables equations to be written very directly in a high level computer language. Graph plotting and tabulating output routines are included. DYNAMO 2 (a new version) allows FORTRAN statements to be included. Special DYNAMO functions include delay functions of various types, test functions such as steps and pulses, random noise generation and the ability to input functional relationships in the form of tabulated numbers. A time-sharing version of DYNAMO is also available.

Other systems modelling languages exist, for example, CSMP 360,(6) which have certain advantages.(7) The virtue of DYNAMO is the directness with which it relates to the theoretical structure just described.

The application of industrial dynamics

THE basic methodology of industrial dynamics is not hard to learn, but its successful application is less straightforward. All modelling is and must be an art in which judgment and experience play a large part. In industrial dynamic modelling, the necessary judgment is perhaps only rarely combined with the analytic ability also needed. Furthermore, it is seldom, if ever, easy to construct tests of validity for proposed models if they represent highly complex situations.

TABLE 2—OPPOSING VIEWS OF MATHEMATICAL MODEL-BUILDING

Conventional views	*The industrial dynamics position*
1. **Data**	
The main obstacle to the development of useful models of industrial operations is the lack of adequate data. The first step must be extensive collection of statistical data. Analysis of simple relationships between measurable variables should be the starting point of model building.	We usually start already equipped with enough descriptive information to begin the construction of a highly useful model. A model should come first. The model builder is more in danger from being insensitive to and unperceptive of important variables than from lack of information, once the variables have been exposed and defined.
2. **Validity**	
Objective tests must be available to establish the degree of correspondence between the model and the real world and the levels of significance of its hypotheses.	Model validity is a relative matter. The usefulness of a mathematical simulation model should be judged in comparison with the mental image or other abstract model that would be used instead.
3. **Prediction**	
The aim of model building and the criterion for its success and scientific status is its ability to predict the future states of the system.	System models should predict and reproduce the behaviour character of a system, not specific events or particular unique sections of actual system time history.
4. **Solutions**	
We should aim at analytic solution of models and this implies, in most cases, limitation to linear systems with constant parameters.	Most dynamic behaviour in social systems can be represented only by models that are non-linear and so complex that analytical mathematical solutions are impossible. For such systems only the simulation process using step-by-step numerical solution is available.
5. **Generality**	
The level/rate structure enshrined in the DYNAMO compiler is of limited applicability. It describes a small sub-set of industrial situations.	The structure will come to be recognised as having simple elegance, universality and a fundamental character common to a very broad range of systems running from physical devices through medicine and psychology to social ecological systems.

These factors, together with the large claims often made by industrial dynamics practitioners for the universality of their results have made the topic the subject of some controversy among management scientists. Table 2 sets out in extreme form some of the differences that emerge.

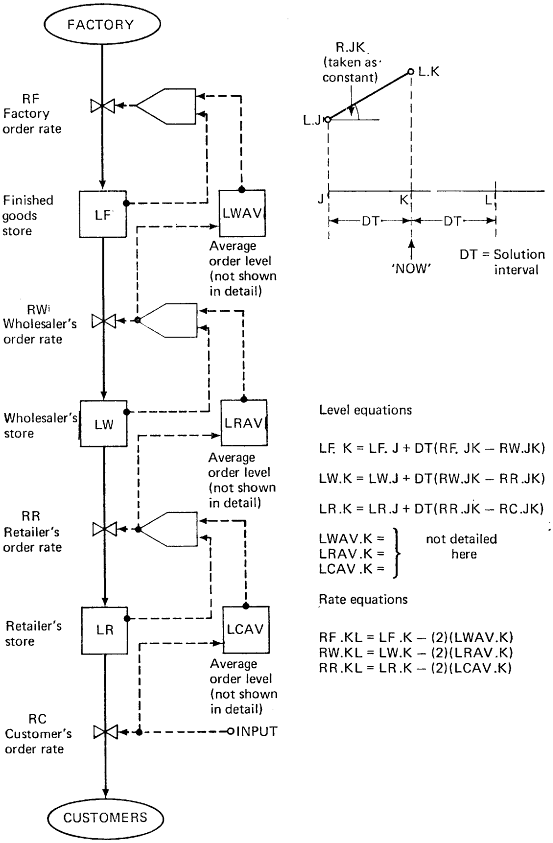

Level equations

LF. K = LF. J + DT(RF. JK – RW.JK)

LW.K = LW.J + DT(RW.JK – RR.JK)

LR.K = LR.J + DT(RR.JK – RC.JK)

LWAV.K =
LRAV.K = } not detailed here
LCAV.K =

Rate equations

RF.KL = LF.K – (2)(LWAV.K)
RW.KL = LW.K – (2)(LRAV.K)
RR.KL = LR.K – (2)(LCAV.K)

Fig. 3—Set of equations (simplified) for 3 stock example

Misunderstanding of the ways in which industrial dynamics is intended to be used accounts for much of the criticism and some of the difficulties in application.

An industrial dynamics study starts with a problem resulting from a system showing aberrant behaviour of some kind, a resistance to change perhaps or a persistent oscillation not obviously related to system inputs. It then attempts to show how the behaviour results from a defect in system structure. Finally, by experimentation on the simulation model, it suggests changes in system policy or structure that remove the self-generated disturbances.

The emphasis throughout is on organisational design or redesign made effective by an understanding of the system dynamics, particularly its sensitive points of leverage.

The work done at M.I.T. and elsewhere has covered many fields extending from the early work on production-distribution systems to broader studies of corporate growth and market dynamics most recently to large-scale research on social systems.

The lessons from this work confirm that—

1. Systems behaviour is frequently remarkably insensitive to parameter change sometimes even by factors of up to five.
2. There are a few sensitive points where small changes can produce radical changes in system behaviour.
3. Prediction of behaviour is beyond unaided intuition and requires a disciplined structural investigation using computer simulation techniques.

It seems likely that only by improving our understanding of system structure in the way these conclusions suggest will we resolve the classic problem of the control engineer, now shared by practising industrial managers. The problem is—How can we create organised systems that are at once stable and responsible to rapidly changing inputs?

There will be no simple solutions, but we may learn how to produce complicated ones.

An example from the paper industry

THESE ideas can be made more concrete by sketching a possible ID approach to a practical problem, which could emerge in the paper industry. It has to be made clear that what follows is in no sense the result of an actual study; however, I have taken the circumstances of an actual situation so as to make the material plausible. What I have chosen to do is to take the production-distribution system of a papermaking company and to examine the possible implications of a decision to change it in a fairly radical way by the introduction of a central warehousing function.

Let me start by listing five points I want to illustrate by means of this example—

1. The aspects of interest in a dynamic study differ from those considered in the more usual static one.
2. There are differences as well as similarities in the characteristics of industrial management systems and engineering systems.

3. Understanding the characteristics of systems of this kind often requires bold use of 'soft' data. (The further point that this is much less hazardous and justifiable than might be thought cannot be demonstrated without an actual study.)
4. The definition of the system's boundary is crucial in modelling.
5. Light can be thrown on the transient effects of changing company situations that tend to preoccupy senior management are often perceived in terms of crisis, usually receive no systematic attention and are, in fact, closely related to problems well studied by control engineers.

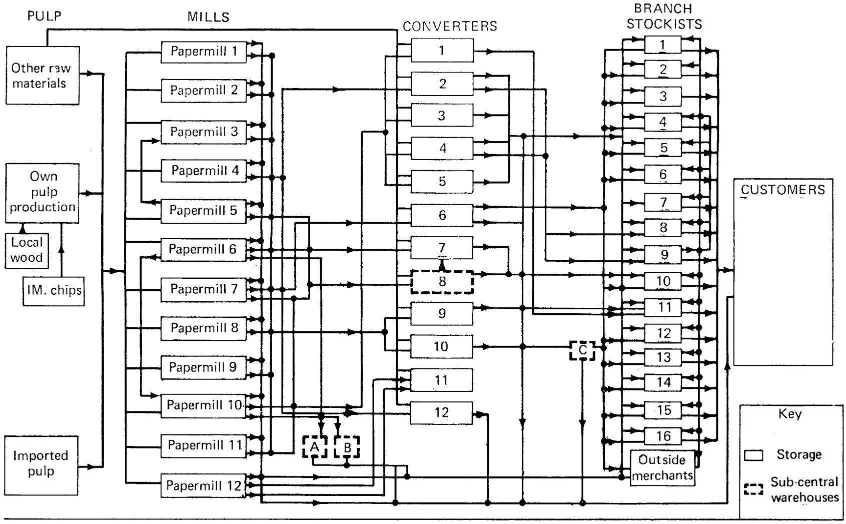

Fig. 4—An existing production-distribution system

Fig. 4 shows in schematic form the configuration of an existing production-distribution system. Typically, a chain in this system works as follows. Papers of a particular grade, but of different weights and finishes are produced at a mill in a repetitive sequence of runs. A stock of the various grades of paper is held in reel form to feed the converting process and a stock of finished paper products allows the response to orders from the branches to be independent of the production cycle. Large customer orders can be slotted directly into the production cycle. Neither of the two stocks is designed to smooth the production process as stock reorders are calculated to replenish the stock to a predetermined level on each cycle, based on an average of past sales. This system allows considerable flexibility in that it can react very quickly to changes in demand pattern, order mix or customer priority. Its disadvantages are that it is unlikely to provide the theoretical savings possible with a few centrally sited warehouses that would allow optimum use of bulk transport and economic order quantities. The use of central warehousing might also be expected to result in reduced delivery times. We shall return to this in a moment.

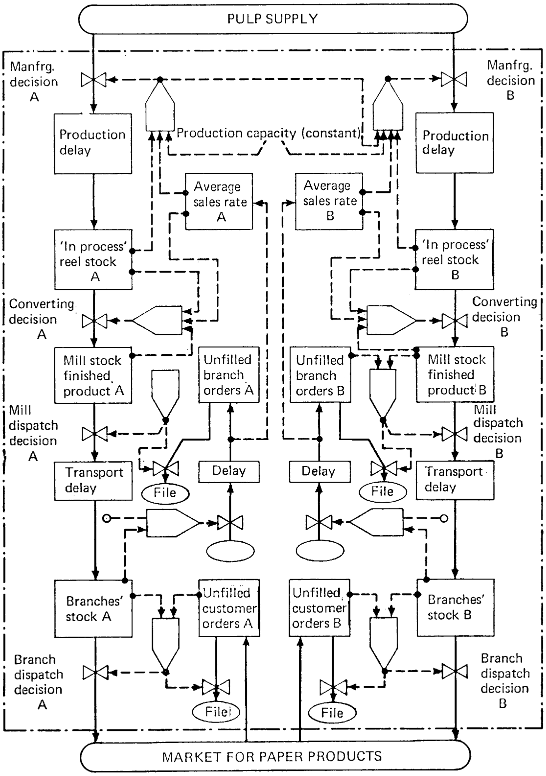

Fig. 5—Simplified flow diagram for production-distribution system (two paper product groups only)

Firstly, let us look at the way in which industrial dynamics could represent such a system. Fig. 5 is a flow diagram based on the model developed by Forrester⁽¹⁾ drawn for just two paper products. It shows how the rate of flow of product from pulp to customer is determined by decisions based on information about the state of affairs in other parts of the system. It allows the introduction of delays at various stages in the process, including lags in information flow. The equations that convert information into rate decisions define the policy structure of the system.

Notice two things—

1. We have drawn systems boundaries at particular points. Pulp buying and storing decisions, plant capacity decisions and the market are all outside this boundary.
2. We have limited the model to two product lines.

Decisions of this type are not arbitrary, but they can be made sensibly only if there is a clearly defined purpose (usually the solution of a specific problem) associated with the model. The point that there is no such thing as *the* model of the system is an obvious one, but nevertheless often forgotten.

For example, we might investigate the following. If the capacity of the papermachines is almost fully utilised, preference is likely to be given in the manufacturing decision to those grades of paper that allow higher tonnage rates of production. The effect of this would be to increase average delivery delays on the less favoured products. What system behaviour might result from this? The model builder would clearly have to include some representation of the market in order to show how order rate would react to the longer delivery time. He might want to include decisions on the increase of papermaking capacity. He would certainly have to consider how many different lines he should include in his model and an effective representation of the priority production decision. This could lead perhaps to studies of a long-term realignment of customers following fluctuations of sales linked with the state of the economy.

The model that would result from this would be totally different from the one limited as in Fig. 5.

Let us return to this simpler model. Fig. 6 shows some typical order data. A high level of noise is imposed on long-term trends. A good deal of this noise may be the result of the discrete nature of branch ordering, perhaps times to coincide with known production schedules. Three points can be made here—

1. An important difference between managed industrial systems and many physical systems is that there is in the latter often a fairly clear separation of signal frequencies, noise frequencies and the natural frequencies of the system, whereas in management systems this is only rarely so.
2. In industrial dynamics, attention is focused on the continuous nature of the systems it studies. This is not to say either that the discrete characteristics of systems are unimportant or that they cannot be repeated in an industrial dynamics formulation, only that most of the interesting dynamic behaviour can be modelled with continuous approximations to the system elements.

3. In industrial dynamics studies, systems behaviour generated in noise-free decision processes is emphasised. This is in contrast with event-based simulation methods, which try to generate the stochastic behaviour of the system. Yet noise signals are often used to test the sensitivity of the system to random inputs.

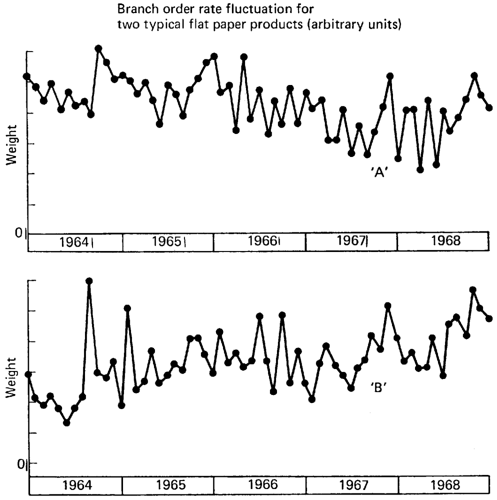
Branch order rate fluctuation for two typical flat paper products (arbitrary units)

Fig. 6—Some typical order rate patterns

It might appear with demand data of this kind that the dynamic response characteristics of the system were of minor importance. Certainly, in industries subject to rapid swings in fashion or frequent production innovation, inadequate dynamic response is more immediately noticeable. Yet no industry —however static it may appear—can neglect its dynamic behaviour. There are few industries, for example, unaffected by the external state of the economy. Devaluation, lowering of trade barriers, changing cost of capital, changing levels of consumer expenditure can in their various ways set off chains of reactions that, if inappropriate, can be unnecessarily expensive. The need to detect change in the presence of noise entails an averaging process that itself introduces a lag.

The other source of dynamic effects nearly always present is competition. One of the ways in which companies compete is by offering better customer service and, in particular, shorter lead-times on delivery. In our example, we can take it that pressures to improve delivery provide the impetus for the central warehouse proposal.

The 'best' configuration of production and storage locations can be worked out straightforwardly, which is not to say easily, by operational research techniques fairly well established. Such a study might show an impressive saving both in cost and in delivery time, but it would be very likely to be a static study rather than a dynamic one. It would thus say nothing about the effect of change in the new system, whether this change be an external effect in the economic or market environment or the internal impact of the new system itself.

To end this paper, a short sketch is given of how industrial dynamics might study these aspects of the proposal.

We noted in the account of the existing production-distribution system that the production decisions were made in direct contact with incoming orders, in some cases from customers or at least from the branch level. If central warehousing were introduced, this contact would be lost and production control would infer correct production rates and mixes from the data available to it. This would require the setting up of a new information system. Alternative systems could be studied using a development of the model of Fig. 5, modified to include both the central warehouse and the new information handling procedures. The tendency of these systems to amplify noise inputs in the way we have referred to would be examined using a model of this type. One might look also at the effect of using information on other variables such as delivery delay on system operation. If it were favourable, new measurements might be worthwhile making.

We could turn next to the study of possible reactions to the warehouse introduction itself. One of the characteristics of large systems we have remarked on is their ability to resist change. This may not be the result of some natural conservatism among the people running them, but be caused by the hidden closed loop structure of the system. For example—

1. There could be a change in customer ordering patterns corresponding to a lower stockholding by them when they perceived the shorter lead-time on delivery.
2. Failure to signal a clear change in goals corresponding to the potential of the new system may result in self-regulating actions in related areas that nullify much of the value of the change.

Detailed examination of effects of this kind could be conducted by a model outlined in Fig. 7, based on one developed in detail in Forrester's paper.(8)

We can see from this feedback loop representation how central a part delivery delay plays in the operation of the total system. Firstly, it decreased the effectiveness of sales effort, hence decreases orders. Secondly, if it differs from the company norms' for delivery, delay pressures will be created for additional plant capacity. If we add the positive feedback loop that tends to increase the sales force as sales rise (or decrease the force if the sales fall), we

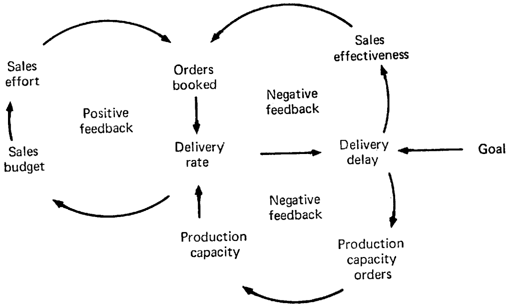

Fig. 7—Three coupled feedback loops

have three interacting loops that can exhibit a variety of behaviour difficult to predict without simulation. The effect of differential delivery delays on products of varying profitability provides a further dimension of complexity.

All we have been able to do is to sketch an outline of certain speculations about system behaviour. Because the actual study has not been done, they have been based on models developed in other situations by Forrester and his associates. In the nature of things, the complexity of system behaviour cannot be described verbally in a few sentences and this has not been tried. Indeed, to succeed would have contradicted the central theme of this paper.

The purpose has been simply to stimulate control engineers and others to widen the boundaries of the feedback systems they study to include factors beyond physical ones. The work will be difficult, but the pay-off may be high. More should make the attempt.

References

1. Forrester, J. W., *Industrial Dynamics* (The M.I.T. Press, Cambridge, Mass., 1961)
2. Roberts, E. B., *The Dynamics of Research and Development* (Harper and Row, New York, 1964)
3. Forrester, J. W., *Urban Dynamics* (The M.I.T. Press, Cambridge, Mass., 1969)
4. Street, P. R. and Dent, J. B., 'Industrial dynamics and an approach to its use in farm management research': *The Farm Economist*, 1968, **11** (8), 345–353
5. Pugh, A. L., III, *DYNAMO Users Manual* (Second edition, The M.I.T. Press, Cambridge, Mass., 1963)
6. Brennen, R. D. and Silberberg, M. Y., 'Continuous systems modelling programs': *IBM Systems J.*, 1967, **6** (4), 242–266
7. Charlton, P. J. and Street, P. R., 'Some general problems involved in the modelling of economic systems on a digital computer' (A.R.C. Symposium on the Use of Models in Agricultural and Biological Research, February 1969)
8. Forrester, J. W., 'Market growth as influenced by capital investment'· *Ind. Management Rev.*, 1968, **9** (2), 83–105

Discussion

Dr J. Grant I have had occasion to make a number of market studies, mostly in connection with my work in developing countries, where the problems are rather different from those outlined by Mr Morgan. Usually, there is no existing papermaking industry and one must try to ascertain in advance what sort of paper the one-machine mill that is going up should make from the materials available. In the course of years, I have developed a technique of my own, making use of such elementary statistical expertise as I possess and, in general, it has not worked out too badly. Often, I do this with some trepidation, partly because the information I receive may be unreliable and partly because the installation of just one converting factory can completely upset my deductions. If I decide, for example, on a general purpose mill for making cheap writings and printings, then, if a cement factory starts up, the demand for sack paper could be far in excess of the total other paper and board requirements of the country; a kraft mill would then obviously be the best type of mill to choose. Is this ad hoc approach really the best one in the circumstances or would it be better to try to apply Mr Morgan's more sophisticated methods?

On the kind of long-range forecasts referred to, if they become generally known, they can often trigger off a set of circumstances that make the forecasts quite wrong. A classical example occurred in Europe about eight years ago when a well-known international organisation forecast a great increase in woodpulp consumption over production three or four years later. As a result, the woodpulp mills put in much plant to increase production; when the three or four years had elapsed, there was a surplus instead of a deficit. The market has not really settled down after that episode even now. If long-range forecasts of that kind are given too much publicity, therefore, they can become a menace.

Mr N. B. Morgan I agree with your first comment, for I am all for ad hoc methods. I know nothing about developing countries' economies, but, if you take the very simple situation you referred to, then the kind of methods I have been talking about are probably not very valuable. It is much better merely to ask the consumer how many rough paper sacks he wants.

I am rather against long-range forecasting in economics, except in the very broadest sense. One believes a growth rate for this country is going to be 3 per cent per annum, for example, instead of 10 per cent. Long-range forecasts for small sectors of the economy must be very doubtful. The outcome for a particular industry or plant can be affected by other people, quite ignorant of it, probably quite uninterested in it, operating on some quite different sector of the conomy with entirely different objectives. Thus, the Chancellor of the Exchequer may have nothing against motor cars; nevertheless, with other considerations in mind, he can make tax and other changes that throw all forecasts out of line.

Dr R. A. Holm The previous statement on pulp capacity has a direct bearing on the point that was made about the necessity of studying the detailed dynamics of economic systems and, proceeding in this pattern, I think the second paper pointed the way very clearly. How do we logically lay down what we believe to be the particular pattern of action or reaction of a system? Does Mr Anderton have any particular methods that he can recommend to determine the non-linear coefficients? It is obviously impracticable and perhaps impossible to determine what these functions are by simply looking up all the past results of similar actions. How do you, in fact, get the main factors and their relationships?

Mr R. H. Anderton There is no simple answer to that question. You arrive at a model by doing rather painstaking informal investigations like interviewing, by data collection in the various parts of the company that seem to be important and by slowly trying to build up an understanding of what is happening. There is a certain kind of man who will do nothing until he has all the data and will only then proceed to analyse it. Industrial dynamics takes the position that usually there is enough known in a descriptive, verbal way about the operations of the company to make a first attempt at making a model. (There is often confusion here between precision and accuracy. One can study something precisely—that is, in numerical terms—without knowing about it very accurately.) A preliminary model is then constructed and one performs experiments with the model to find out which parts of the model (gain factors, non-linear relationships, etc.) really matter—that is, to which the model is sensitive. Having worked that out, you can then go back to collecting more data in that sensitive area. One of the most remarkable results of Forrester's work is that the behaviour of very large systems seems to be quite insensitive to the values of very many parameters, even through factors of five. What it is usually very sensitive to is the structure of the system, in particular, the information that is available and being used at the various decision points. Some people tend to say that they know nothing about that particular variable, therefore they leave it out. In fact, by doing so, they are putting it in, giving it a rate of zero—and you can usually do much better than that! You may know it is seven (maybe eleven), but certainly not zero. Usually, it is more important to put all major influences in, even roughly, than to put emphasis in the early stages on an elaborate empirical investigation of the data.

Mr H. B. Carter On my way to this conference, I ran into a trades union official and this very problem came up. He accused the profession that I represent and of which I am proud to be a member of not being able to do anything about the wage/price spiral. I am sure we have enough brains in the profession now to be able to predict the result of any of the various actions and control activities that are embarked on generally by governments. He claimed that we need a deal of statistics and all the facts. I said that all we need are sampling techniques and we can then arrive at all the factors, until we come to a very pertinent point. What happens if, as a result of this control activity, the trade unions are told that the way to stop this is to first of all put a freeze on wages. He answered, 'That would be political suicide; we could never accept that.' So the point of this comment is how do you program such a reaction into an appropriate control when 50 per cent of the labour force who are to be directly affected will reject it anyway?

Mr Morgan First of all, there seems to me to be a reasonable chance that a number of ways of controlling the situation exist—indeed, any situation—and reasonable therefore that, of these ways, some of them will have results worse than the complaint for at least some of the population. Many people regard a wage freeze as being worse than having a stop/go cycle.

Mr Anderton There is a very interesting class of control system in which what is controlled is very complicated and is, in fact, much more complicated than the complexity of the controller. The variety of response available to what is controlled is thus greater than the number of control actions available to the controller. This happens very much in national control systems, when the Chancellor of the Exheuqer introduces a new tax system, for example. Then there may be thousands of people who are trying to find ways of evading tax. There are more of them than there are of him. This raises a fundamental point that can perhaps be tackled only by the mathematical theory of games. It is a whole new class of control system.

Mr A. J. Ward May I return to Mr Carter's question. I think, there was some confusion in his comments. These techniques do not generate control strategies for you, but enable one to demonstrate the dynamic behaviour of the system studied.

On the second part of his question, engineers tend to fight very shy of problems that are not physical and are not easily quantifiable such as marketing, selling, promotion and investment. These are often dominated by dynamic phenomena and control systems engineers have just the right background to formulate and analyse them. Even if the numerical accuracy of the answeer is poor, a system study will often show marked differences in the quality of the system's general behaviour for different varieties of control action. Here is a real opportunity for engineers to make a contribution to the non-technical aspect of the business.

Mr P. E. Wrist I would like to say to Mr Anderton that I think the techniques he is discussing are probably very effective in periods of stability or in periods of continuing expansion. They tend to break down during transition periods. I would like to mention three examples.

One is the departure from normal behaviour during periods of national uncertainty about the future economic climate or a loss of public confidence. During these periods, for example, inventories may either be run down to a very low and dangerous level in anticipating a recession or run very high in anticipating a shortage or a strike. Most of the models of which I have been aware are not adequately able to predict and incorporate these things in time. They can usually analyse what happened after the event, but they often cause misleading projections of demand for the short term and a distortion of demand over the long term. If we followed the predictions, we might conclude that a new expansion of capacity was required immediately only to find a few months later that the assumptions on which it was based have disappeared. Looking back, we can analyse why we were misled, but the model did not prevent us from being misled.

Another factor that can distort demand projections is a changed ordering pattern that happens during times of shortages. We have found that during these periods there is an excessive amount of double and triple ordering. As soon as shortages begin the ease, these duplicate orders vanish overnight and arrears that previously appeared to be as firm vanish in a very short time.

Another effect that makes economic models of limited value is their inability to predict the unusual. Events that happen only once would never be incorporated into the model, because of no previous experience and the probability that they will not happen again. By way of example, if we look at the book industry in U.S.A., 1967 was a year of very rapid growth and forecasts for continued growth were made. Many plans were made on that assumption. The following year, the expansion did not occur and, on investigation, it was found that the government towards the end of the previous fiscal year had a large sum of money left for education and the only thing that people could buy in a short time was school books. As a result, in the summer of 1967, an excess of school books was bought. In 1968, the stock was drawn on and fewer new books were bought. This kind of action will probably not happen again for a long time; when it did, it had a dramatic effect on a system such as you described.

Mr Anderton I agree with most of what you say. No technique can prevent you making mistakes in prediction, but if you have a technique for understanding the mechanism of events, even post hoc, then you will at least not make the same mistake twice. With an increasing understanding of the way the systems are working, I think we will be able to anticipate events more frequently.