The Planning of New Chemical Processes*

By Dr. D. H. ALLEN

Department of Chemical Engineering, University of Nottingham (England)

Summary

This paper describes various factors which affect new chemical processes and how they are being investigated to try to improve methods of planning and control. The factors considered are the economic aspects of projects, the accuracy of forecasts and allowing for uncertainty, the selection of research projects and the allocation of resources to them, and the accumulation of information as a project proceeds.

Introduction

At the Chemical Engineering Department of the University of Nottingham we are carrying out research into the selection, planning and control of new projects in the chemical and allied industries, particularly those involving research and development. Our main interests are in methods of presenting, analysing and processing information used to make decisions about projects, in the way decisions are made, and in the consequences of these decisions. The approach we use is essentially a practical one, based on the study and analysis of actual projects in industry—both past and present, successful and unsuccessful, and it includes the detailed examination of their scientific, technical, commercial and financial features. You may wonder why this work is being done in a University Department of Chemical Engineering, since the research is not primarily into the chemical and physical aspects of processes, which are the more conventional fields for such establishments. Our interpretation of chemical engineering is wider than this. Anything which affects the success of a chemical process we consider can legitimately be included. The scientific, engineering, economic and management aspects of a project are closely inter-related and need to be viewed as contributing to the project as a whole. The chemical engineer is well-fitted to be able to do this.

Economic assessment

During the planning of a new project several important questions may have to be answered, for example:

- is a new idea worth developing?
- if resources are limited, which projects should be given priority?

* Paper presented to a meeting of the Basel Chemical Industry on 6th January 1969.

- should a project be stopped now to cut possible losses?
- should the direction or objectives of a project be changed?

Such decisions require some sort of assessment of the situation. A project assessment or evaluation consists of balancing benefits to be expected against the resources needed to achieve them. It may be done intuitively by an experienced management, but this is becoming more difficult and less effective as situations become more complex, and it is usually advantageous to quantify the assessment as far as possible. In order to compare eventual benefits with the use of resources quantitatively, they must be measured in the same units, and the universal ones are those of money. The economic evaluation of a project is carried out in terms of the costs of bringing it to fruition and the income or savings which will result.

In an economic evaluation of a project in its early stages, it can be unwise to consider only research and development without reference to other activities needed for its successful exploitation. Perfecting a piece of technology may be of little value unless financial and other resources are available at the right time to exploit it, and so reap the benefits. We have an example of this in one of the actual industrial projects we have studied. A small company did some good research and successfully developed a new petrochemical process which opened up a new field of manufacture. Rather than licence the process to the industry, the company decided to go into production itself, but it only built a small plant because of limited capital and lack of experience. Because of its small size the plant had high unit operating costs and did not make a satisfactory profit. The company was soon in financial difficulties with its backers, and was eventually taken over by a larger company who immediately expanded the plant and made the new process very profitable.

The economic aspects of a project can be presented by means of a cumulative cash flow diagram as shown in Fig. 1. The vertical scale is the cumulative cash position, either negative when the project is in debt, or positive when it shows an overall profit. The horizontal time scale is measured from the start of the project. AB represents preliminary expenditure on research and
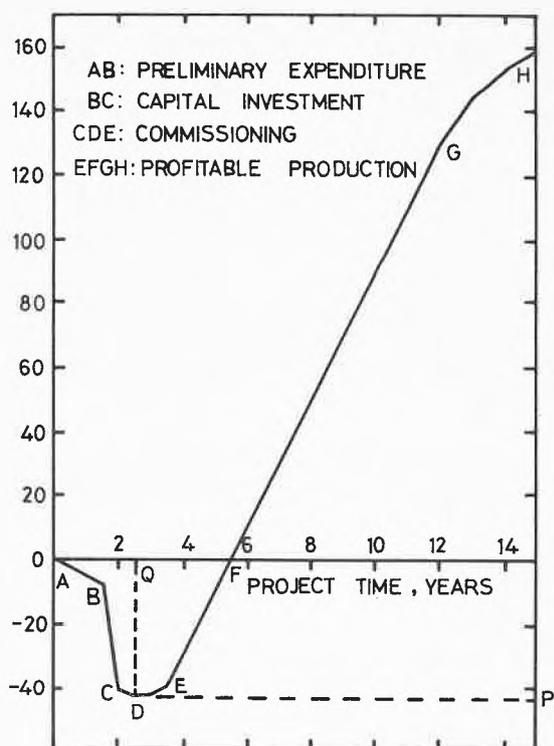
Fig. 1. Typical cumulative cash flow diagram for a project

valuable material in short supply. After some preliminary laboratory work production was started using a pilot plant on a small scale, and it paid for itself before the investment in a larger plant was made. The steep rising part of the curve shows that it has been a very successful and profitable project. These diagrams use a normalized cash flow scale in which cash positions are given relative to the maximum cumulative expenditure. This assists in the comparison of different projects.

The financial attractiveness of a project is a function of the shape of its cash flow diagram. Several different functions are in use. Each function, or method of economic evaluation, identifies important features of the cash flow pattern and converts them into a simple numerical measure. I am not going to discuss in detail the functions which can be used, but in general they fall into three categories in which the units of the functions are those of money, time or a rate of return. Two examples from each category are given in Table 1.

Table 1. Methods of economic evaluation

Measure	Units	Examples
Money	£, $, Francs, etc.	1. Net cash position after n years 2. Present worth
Time	Years	1. Payback time 2. EMIP
Rate of return	Percent per year	1. Return on investment 2. DCF return

development and this is followed by *BC*, the capital investment in plant. *CDE* is when the plant is commissioned and the curve turns upwards as a net income is obtained from product sales or cost savings. Eventually with a successful project the cash position will become positive and continue to rise.

Examples of the many cash flow diagrams of actual projects which we have studied are given in Fig. 2 and Fig. 3. Hector is the code name for a new process to recover vapour of a valuable solvent from the effluent gases of a production plant on a large scale. Eight years elapsed from inception to the successful start-up of the first major plant. Hermes is the code name for a project to convert an almost worthless by-product into a
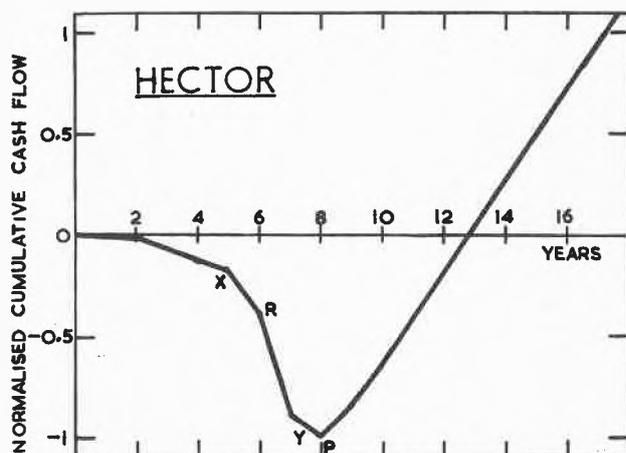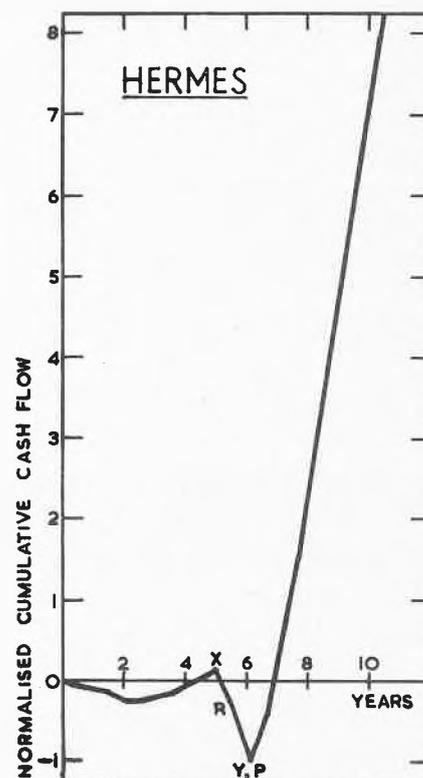

Fig. 2. Cash flow diagram of Project Hector

Fig. 3. Cash flow diagram of Project Hermes

Table 2. Project present value and DCF return

C_t = cash flow in year t
n = project life
i = discount rate
C = cash flow as a function of t

	Present value V	DCF return I
Cash flow in yearly increments	$V = \sum\limits_{t=0}^{t=n} \dfrac{C_t}{(1+i)^t}$	$\sum\limits_{t=0}^{t=n} \dfrac{C_t}{(1+I)^t} = 0$
Continuous cash flow	$V = \int\limits_0^n C \cdot e^{-it}\, dt$	$\int\limits_0^n C \cdot e^{-It}\, dt = 0$

Present value and DCF return both take account of the so-called time value of money by discounting later cash flows back to start of a project, according to the equations shown in Table 2. The first equation gives the present value V of a project with n yearly cash flows, C_t (positive and negative), discounted at a rate i percent per year. The second equation gives the continuous form of the same relationship. For DCF return, the value of i is found which makes V zero. Present value and DCF return are illustrated in Fig. 4. Curve 1 is the original cash flow curve with no discounting. Curve 2 represents the present value of these cash flows discounted at 10% per year, and JK is the present worth of the whole pro-
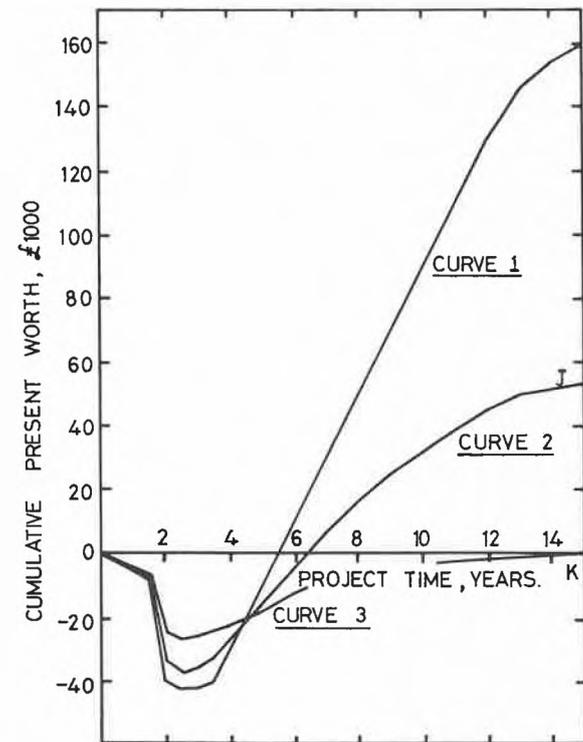

Fig. 4. Cumulative present value diagram for a project

Curve 1: No discounting
Curve 2: Discounting at 10%. Present value = JK
Curve 3: Discounting at DCF rate of 29%. Present value = 0

ject. For this project, if the discount rate is increased to 29% per year, curve 3 results, in which the present value is zero, and so the DCF return for this project is 29%. Another way of looking at DCF return is to regard it as the interest rate at which money could be borrowed to finance a project and just breakeven at the end, after repaying capital and interest charges. Curve 2 of Fig. 5 indicates the cumulative cash position at any time for the project, including interest payments at the DCF rate of return.
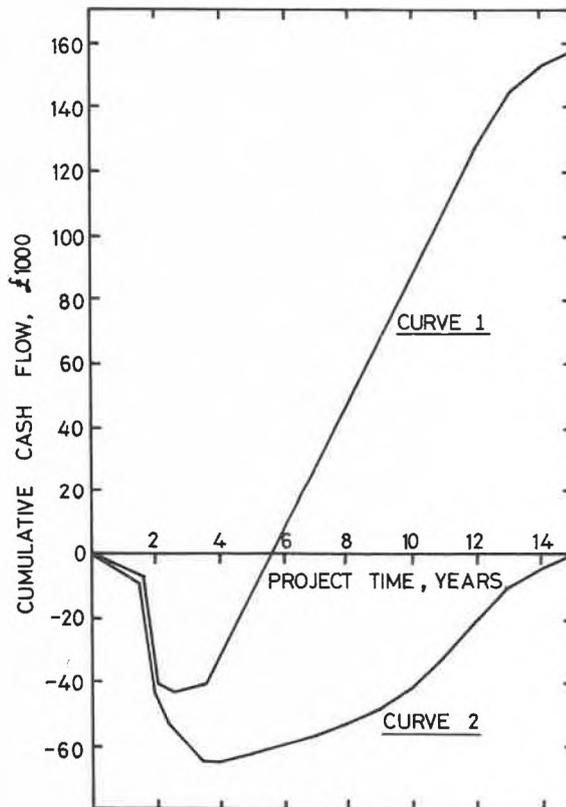

Fig. 5. Cumulative cash flow for the same project, including interest charges at DCF rate (29%)

Curve 1: Original cash flow with no discounting
Curve 2: Cash flow including interest charges

A new function which we have developed as a result of our studies is called EMIP. This is a short-term function which has been found to correlate very well with a project's profitability. As shown in Fig. 6, EMIP is the area enclosed by the cash flow curve below the zero on the cash scale, divided by the maximum cumulative cash debt. It is equivalent to one side of a rectangle of the same area, whose other side is the maximum debt. It thus represents the equivalent period during which the maximum debt would be outstanding if it were all incurred at the same instant and all repaid at one later instant; hence the full name Equivalent Maximum Investment Period. The shorter the EMIP the more attractive a project is. EMIP is a function which takes
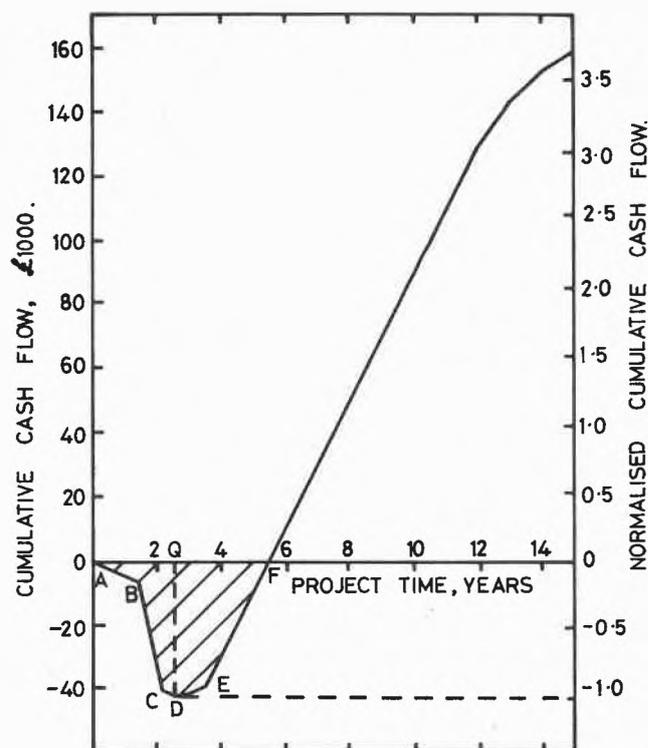
Fig. 6. Cumulative cash flow diagram for a project showing measurement of EMIP

$$\text{EMIP} = \frac{\text{Area } (ABCDEFQA)}{QD} = 2.7 \text{ years}$$

On normalized cash scale, $QD = 1$ cash flow unit

account of the pattern of cash flow as far as the break-even point and has been found to be useful in the preliminary screening of projects when the long-term position may be unclear and the later cash flows are very speculative.

Forecasting problems in new projects

When a new project is evaluated in its early stages the future cash flow is of course based on forecasts of costs, market conditions and timing. These in turn reflect the technical and other problems which have to be solved. We are also studying projects to see how forecasts about them compare with subsequent events and to find reasons for differences. We hope that some general guidelines for forecasting procedures and information accumulation during the research and development stages of a project will emerge from this work.

I can illustrate this with the example of a project we have studied which lost a lot of money. The process was a fairly simple one, chemically speaking. It was to make a particular organic compound on a large scale in a continuous process plant. The reaction required ammonia and initial laboratory research was carried out on two alternative processes, one using a 400 litre vessel and aqueous ammonia, and the other using a one litre autoclave with anhydrous ammonia. It was decided at

an early stage to concentrate further effort on the anhydrous ammonia process because corrosion problems would be less severe and the plant would be smaller, although not necessarily cheaper. The company's interpretation of the commercial situation was that production must start quickly to corner a new market before home competition and foreign imports gained a foothold. Because of this, the continuous pilot plant stage of development was omitted and a production plant was designed on data from the one litre batch laboratory equipment. The plant was built and commissioned quickly but soon ran into trouble in operation. The reactor tubes became blocked every few days and this meant frequent shut-down for cleaning them was needed. To find the cause, this full-scale plant had to be operated experimentally. It was identified as a polymeric material produced by the reaction of two by-products, one of which was recycled with the excess ammonia and so accumulated in the system. Its concentration was too low to be detected in the batch laboratory equipment, but it would undoubtedly have been found in a continuous pilot plant as it accumulated. During those delays the price of the imported product came down. Proposals were made to solve the technical problem by major modifications to the reactor system. With this additional expenditure and the reduction in the market price, economic operation could only be possible if cheaper raw materials could be found. This was not possible and the plant and the project were abandoned. The cash flow for this project in terms of forecasts made and what actually happened is summarised in Fig. 7. It can be seen that forecasts were always very optimistic compared
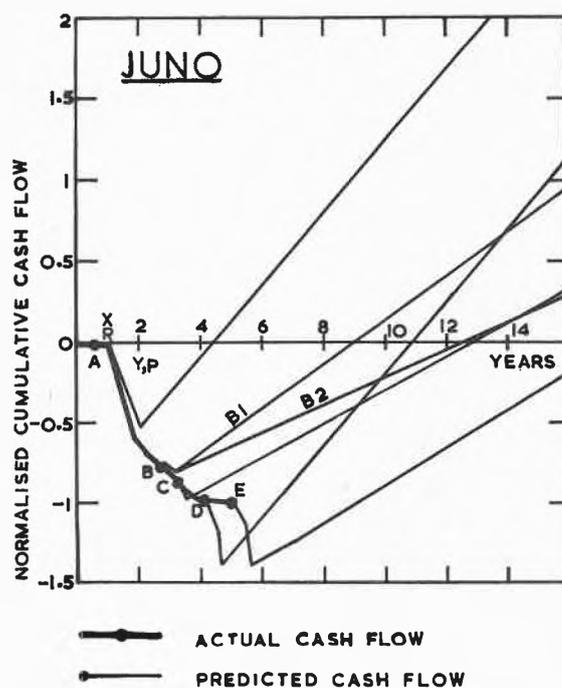

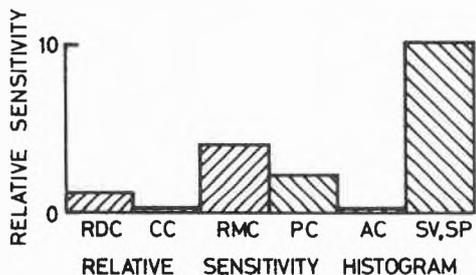Fig. 7. Predicted and actual cash flow for Project Juno

Fig. 8. Example of the sensitivity of profitability to proportional errors in component forecasts for a project

R D C   Research and development costs
C C   Capital costs
R M C   Raw material costs
P C   Production costs
S V, S P   Sales volume, selling price

with what actually happened. Initially the company did not appreciate the risk it was taking in not running a continuous pilot plant before designing the production plant. On the commercial side the company failed to appreciate the implications of the trend in the international market for the product. The failure of the project was due to the combination and interaction of these technical and commercial mistakes.

Errors in some types of forecasts are of greater significance than comparable errors in others. The relative sensitivity of a project's profitability to errors in individual forecasts can be found by analysing the predicted cash flow diagram. As an example we can consider the forecast sensitivity histogram in Fig. 8 for a project which was to make a new herbicidal product. This shows the relative effects on profitability of the same percentage errors in different forecasts. As can be appreciated from the cash flow diagram, the forecasts of sales volume and selling price are always the most significant, as together they represent the income of the project, and this must always be larger than all the cost forecasts put together. For this project, raw material costs were next most important, followed by production costs and research and development costs. Since the new product in this case utilised an existing plant, the capital investment was low. This type of analysis early in a project identifies the more important forecasts and information which is needed, so that the project can be planned accordingly.

We are also gathering data on the actual accuracy of forecasts made at different times in various projects. This work has led us to some new ideas in forecasting which are being tried out. Typical forecast accuracies are given in Fig. 9. They are for a project for the development of a process to make a new compound used in the manufacture of detergents. In each diagram of Fig. 9 the vertical error scale is relative to the actual value A and the horizontal time scale is also a relative one, from 0 at the start of research to 1.0 at the start of production. The crosses represent forecasts made at particular times and the circle the actual result. The crosses are joined in steps since a forecast usually continues to be used unchanged until a new one is made. It can be seen that the accuracy of forecasts varies, some being pessimistic and some optimistic. In particular, forecasts of product sales volume are very optimistic. This is a general conclusion from the study of several projects and has to be considered in association with the fact that, as we
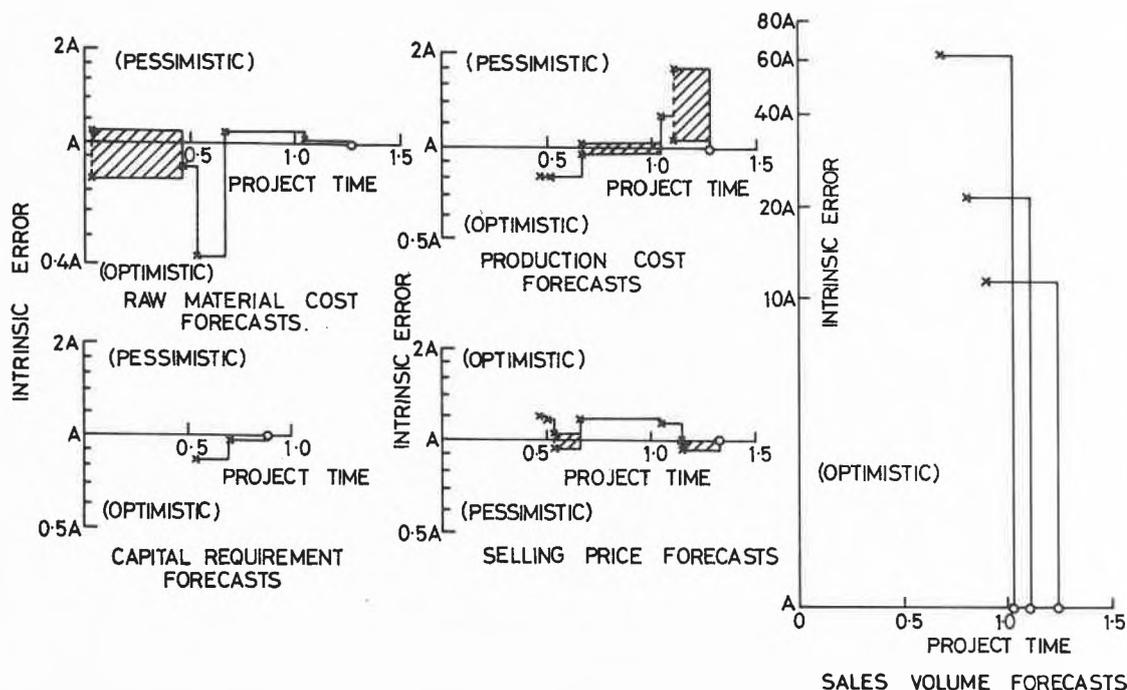

Fig. 9. Example of data on the accuracy forecasts for a project

have seen, sales forecasts are the most critical in an economic evaluation of a project.

Methods of allowing for uncertainty in project forecasts

A simple method of allowing for risks resulting from uncertain forecasts is to reduce the financial attractiveness of a project by a factor which increases as uncertainty increases. This factor is called the probability of success and is a subjective assessment of the probability that the otherwise predicted profit will be achieved. Sometimes the factor is split up into other factors, such as the probability of technological success and then, given this, the probability of commercial success.

A more sophisticated method is to generate probability distributions for individual forecasts and to combine them into a probability distribution for the whole evaluation. Mathematically this is quite complex and the assistance of a computer is needed. Such methods of treating uncertainty were originally developed for statistical forecasting and repetitive investment decisions. Their purpose is to provide help in making decisions by assessing uncertainty and risk. Their value depends on the quality of the information used and the way it is processed and presented. These more sophisticated methods of dealing with uncertainty usually require more detailed and specific information as input, and their usefulness is limited by the quality of the information available. In some situations, including the planning of new chemical processes, the use of such methods on poor data may give misleading results, and in some cases this has mistakenly brought them into disrepute.

We feel that there is a need for better, simpler methods which can take into account the practical limitations on obtainable data and realistic forecasts. In a completely novel project involving much research and development, it may be unrealistic to make forecasts in terms of probability distributions if there is no basis of past experience for estimating probabilities. Instead we can only make range forecasts without any knowledge of the shape of the distribution. We called these *credibility* forecasts. Their use leads to simpler models which are more easily applied to the type of new projects we are considering.
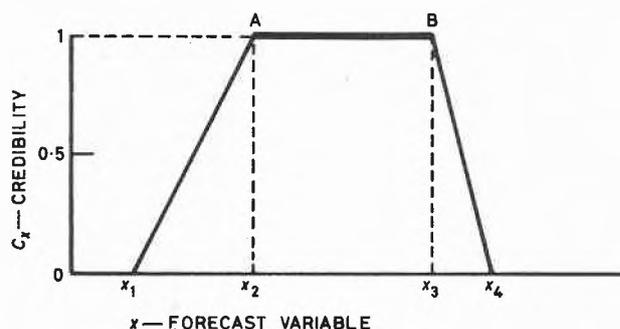

Fig. 10. Forecast range credibility diagram

Fig. 10 illustrates a credibility range forecast. Any forecast value of x between A and B has a credibility of 1, or is fully credible. Values less than x_1 or greater than x_4 have zero credibility, or are incredible. Such a forecast is made by starting with a fully credible value and moving outwards both ways until the forecaster just begins to doubt whether it is still fully credible. This establishes A and B. Similarly the outer limits are found by moving inwards from incredible regions until there is again doubt as to whether the value is still completely incredible. Alternatively, credibility can be interpreted in terms of surprise. If the forecaster would not be at all surprised if a particular value subsequently became the actual one, it is fully credible. If he would be most surprised it is incredible. Intermediate credibilities correspond to intermediate degrees of surprise. Credibility forecasts are always made in such a way that any additional information will narrow the credibility range and not extend it.

Table 3. Comparison of properties of Probability and Credibility

Situation	Probability P	Credibility C
Exhaustive forecast	Distributional $\sum P(i) = 1$	Non-distributional $C_i = 1$ for at least one alternative
Combining alternative predictions	Additive $P(A \text{ or } B) = P(A) + P(B)$	Selective $C(A \text{ or } B) = $ greater of C_A or C_B
Dependent predictions	Multiplicative $P(A \text{ and } B) = P(A) \cdot P(B|A)$	Selective $C(A \text{ and } B) = $ lesser of C_A or C_B

The properties of credibility differ from probability in several important respects. These are summarised in Table 3. In each case the *mathematical* operations of probability are replaced by *selective* operations of credibility. Detailed methods of forecasting, combining forecasts and economic evaluation and interpretation using credibility have been worked out, and we are now testing and improving them by applying them to actual projects.

The selection of groups of new projects and the organisation of a research department

The study of individual projects have now led us to consider the problem of selecting the best group or combination or portfolio of projects where the financial budget, manpower and other resources of an industrial research and development organisation are limited. The simplest way, by ranking projects according to their financial returns, may not make the best use of the resources available. This is shown in a simplified way in Table 4. Here there are four projects—A, B, C and D, each of which has a benefit and uses two resources X and

Table 4. Selection by project ranking

Project	Benefit	Cost of Resource X	Cost of Resource Y	Benefit – Cost
A	10	7	2	1
B	24	6	2	16
C	16	10	1	5
D	25	4	3	18
Resource availability		17	3	

Y. If we rank them according to the net benefits in the last column, the order is D, B, C and A. But as resource Y is very limited, only D can be done and the total net benefit is 18. If instead we select B and C the resource usage is still acceptable and the total net benefit is 21, three units higher than before.

Instead of simple ranking, a method is needed based on mathematical programming which will select the combination of projects which maximises the total benefit within the stipulated resource limitations. A computer model employing the techniques of either linear programming or integer programming can be used. The model in its simplest form consists of the matrix shown in Table 5. The columns of x's represent possible projects and the rows are constraints on projects and resources, with an objective function that maximises the total benefit. The elements of the matrix are the coefficients of the project vectors in each constraint or equation. For example, v_1 is the net benefit expected from carrying out project x_1. The x_1 selection constraint ensures that this project cannot be duplicated, that is, that x_1 cannot be greater than 1. In practice, it will either be 0 or 1, depending on whether the project is selected or not. This is achieved by the particular iterative method of solution employed. The x_1 resource coefficients r_1, s_1 and t_1 indicate the amount of these resources that this project will use if it is selected. The sum of the resource usage of all the projects selected must be less than or equal to the resources available, R, S and T. The solution of this optimisation problem by iterative programming on a computer shows which combination of projects from those included will maximise the total benefit and make the best use of the resources.

At Nottingham we have developed this basic model in cooperation with an industrial organisation so that it simulates the situation in their Research and Development Department and it is being used by their management to provide information to help in selecting and controlling research and development projects and in recruiting new staff. There are at present 27 projects in the model, but several of them are repeated in different versions to give a choice in their timing and the resources they use. There are altogether 40 projects and versions. The resources are expressed in terms of the different types of manpower used, with their associated costs included. In this organisation there are 20 categories of professional and skilled manpower engaged on projects, according to their training, experience and level of responsibility. Some men come into more than one category, for example a particular biochemist is also a competent analytical chemist. In order to include such possible interchange of human resources, additional columns and rows are needed in the matrix. It is also convenient to add extra columns to distinguish between existing staff and new staff in particular categories which will need to be recruited. The complete matrix at present has 87 structural columns and 76 constraint rows.

The objective function in which the total benefits are maximised has been modified to enable the model to take account of risk and uncertainty, and to represent the situation for this particular organisation. Firstly, for each project version, a standard return on any subsequent capital investment needed after research has been deducted from the total benefit, where applicable, to give a discounted present value for the research benefit alone. Secondly, the benefit used is not always the absolute benefit of doing the research. It may be the incremental value of the work being done by the Research Department, compared with an alternative means of achieving the same result, for example, sponsoring research elsewhere, or buying in know-how, or taking out a licence.

The benefits for each project are given in terms of a credibility forecast instead of a single predicted value. The optimum combination of projects and project versions to select have now to be found by a series of computer runs making use of the Monte Carlo procedure of statistical analysis. A point in each project version benefit credibility range is selected independently at random by a random number generator, and these point values are then used as the objective function coefficients to solve the matrix and obtain the optimum project and project version combination for this random objective function. This operation is then repeated with other random objective functions selected from the original credibility ranges. In this way a frequency distribution of different optimum combinations is obtained (Fig. 11),

Table 5. Linear Programming Matrix for Project Selection

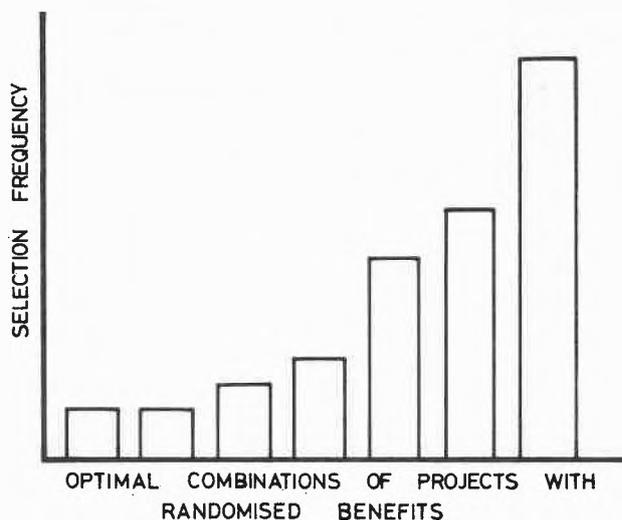|	Projects x_1 x_2 x_3 ... x_n	
Objective function	v_1 v_2 v_3 ... v_n	$=$ maximum
Selection constraints	1	≤ 1
	1	≤ 1
	1	≤ 1
	1	≤ 1
Resource contraints	r_1 r_2 r_3 ... r_n	$\leq R$
	s_1 s_2 s_3 ... s_n	$\leq S$
	t_1 t_2 t_3 ... t_n	$\leq T$

Fig. 11. Project portfolio selection frequency

and this eventually stabilises and shows that one particular combination of projects and project versions appears more frequently than any other. It follows that this most-frequently selected combination of projects and projects versions is the most robust in the face of unpredictable variations in individual benefits within the forecast ranges, and is the best one in the face of the existing uncertainty.

This type of analysis is well within the scope of a medium-sized computer. The model can easily be used to test the effects of changes in any of the data on research programmes, and can be kept up to date as some projects finish and new ones begin. It is providing a useful framework for the detailed organisation of the Research Department on which it is modelled.

Information accumulation patterns in new projects

The last facet of this general field of investigation I wish to mention views the planning of new projects as an information accumulation procedure. Research and

development, market research and other activities needed to launch a new process or a new product are means of obtaining information. The acquisition of information requires time and the use of manpower and other resources. Full information is desirable for planning and controlling a project and for making decisions in order to minimise uncertainty and risk, but spending too much money or time in obtaining it can itself have adverse effects on a project's viability. For any project there should thus be a pattern and level of information accumulation which strikes the best balance between these conflicting requirements and provides the best framework for a project to succeed.

As a first step in investigating how this can be found, we have studied the detailed information accumulation patterns of several actual projects for the development of new chemical products. Individual items of information have been assigned ratings on an index scale, depending on the degree of confidence in the information at the time, or on how thorough, precise and complete the information is. Zero on the index indicates the complete absence of information, while 5 indicates full reliable information. Intermediate numbers indicate intermediate degrees of confidence or reliability. An example of the information patterns found is shown in Fig. 12 for information on the anticipated selling price for the new production of 10 different projects. The relative time scale is from 0 at the start of research to 1 at the start of production. The numbers in boxes represent the information index level at that point in the project. Similar patterns have been prepared for about 40 items of information which affect this type of project.

We are interested in relating information accumulation to the subsequent success of a project. For this purpose we are using three inter-related measures of project success which we call financial, adaptive and execution success. The maximum attainable financial success of a project is set by external factors outside the management's control, such as market conditions. The extent to which this maximum is approached depends on how well the project enables the company to adapt
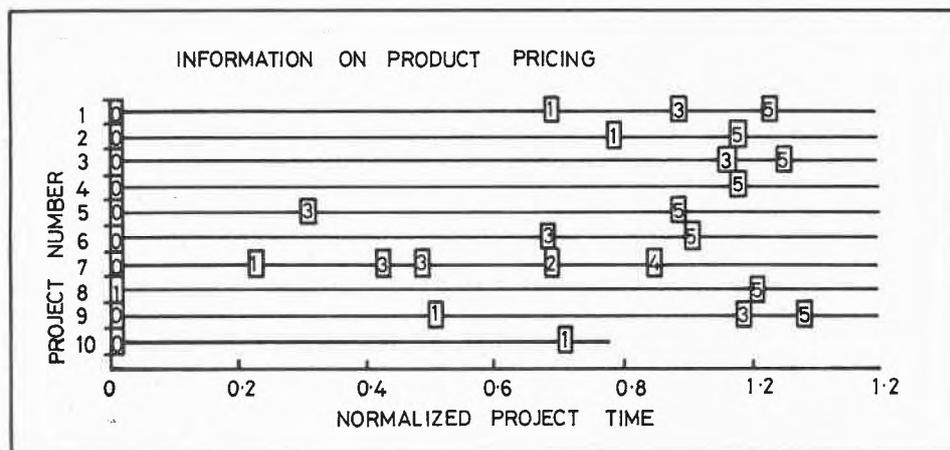

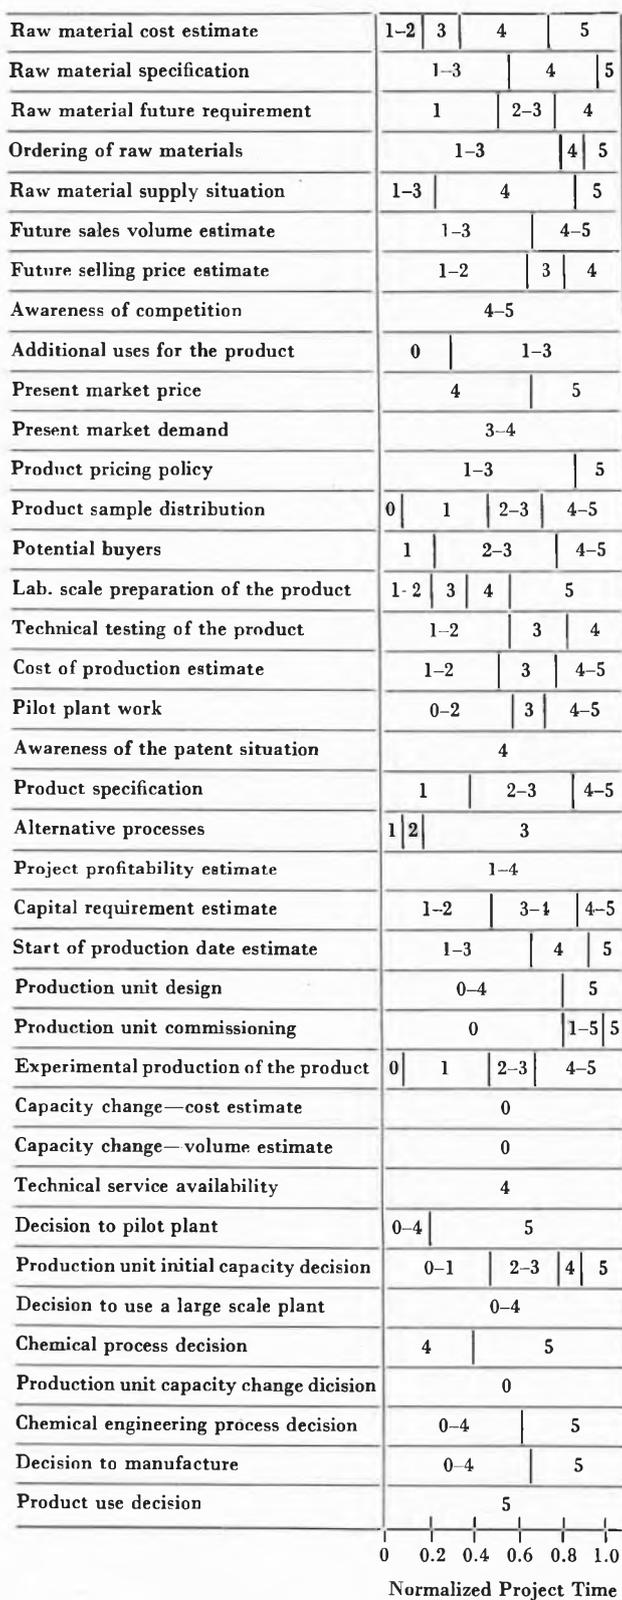Fig. 12. Information accumulation on product pricing for a set of projects

Item	Time bands
Raw material cost estimate	1–2 | 3 | 4 | 5
Raw material specification	1–3 | 4 | 5
Raw material future requirement	1 | 2–3 | 4
Ordering of raw materials	1–3 | 4 | 5
Raw material supply situation	1–3 | 4 | 5
Future sales volume estimate	1–3 | 4–5
Future selling price estimate	1–2 | 3 | 4
Awareness of competition	4–5
Additional uses for the product	0 | 1–3
Present market price	4 | 5
Present market demand	3–4
Product pricing policy	1–3 | 5
Product sample distribution	0 | 1 | 2–3 | 4–5
Potential buyers	1 | 2–3 | 4–5
Lab. scale preparation of the product	1–2 | 3 | 4 | 5
Technical testing of the product	1–2 | 3 | 4
Cost of production estimate	1–2 | 3 | 4–5
Pilot plant work	0–2 | 3 | 4–5
Awareness of the patent situation	4
Product specification	1 | 2–3 | 4–5
Alternative processes	1 | 2 | 3
Project profitability estimate	1–4
Capital requirement estimate	1–2 | 3–4 | 4–5
Start of production date estimate	1–3 | 4 | 5
Production unit design	0–4 | 5
Production unit commissioning	0 | 1–5 | 5
Experimental production of the product	0 | 1 | 2–3 | 4–5
Capacity change—cost estimate	0
Capacity change—volume estimate	0
Technical service availability	4
Decision to pilot plant	0–4 | 5
Production unit initial capacity decision	0–1 | 2–3 | 4 | 5
Decision to use a large scale plant	0–4
Chemical process decision	4 | 5
Production unit capacity change dicision	0
Chemical engineering process decision	0–4 | 5
Decision to manufacture	0–4 | 5
Product use decision	5

0   0.2   0.4   0.6   0.8   1.0

Normalized Project Time

Fig. 13. Optimum information accumulation patterns

0: No information
5: Complete information

Table 6. Success ratings of a set of projects

Project	Success Ratings Adaptive	Execution	Financial
1	4	4	5
2	5	4	3
3	4	3	5
4	2	3	4
5	3	3	2
6	2	2	2
7	1	2	1
8	1	1	1
9	1	1	0
10	(no data)	3 (so far)	(no data)

Ratings   0: very bad   1: bad   2: poor
3: moderate   4: good   5: very good

to changing technology and markets, and also on how well the project is managed or executed. Table 6 shows the actual success ratings for the same ten projects, ranging from 0 (very bad) to 5 (very good). For each project, there is usually good agreement between the different types of success.

All the data on information accumulation and project success is being analysed to see if particular patterns tend to produce successful projects or unsuccessful ones. So far this has only been done by a subjective method of analysis, but means are being devised to make this sort of analysis more quantitative and objective by a network technique. Fig. 13 shows what sort of information accumulation patterns appear, from the data so far examined, to provide the best framework for a successful project. This empirical study has not produced any surprising results, but it has shown that this approach is worth following up. We are now developing the theoretical side by introducing the link between information and success, that is, the decisions made on the basis of information, and on which success depends.

Conclusion

I hope I have been able to give some idea of ways in which this study of new projects is proceeding, and how the results can be useful in planning others. This research at present is principally concerned with the research and development stages of new chemical projects, but similar useful work could also be done in other areas, such as the design, construction and operation of new processes. Chemical engineers and chemists with management responsibilities have to take decisions in all these areas, often on the basis of poor data and uncertain forecasts. Although the ability to do this cannot be properly learned except in the real situation, they can as students be equipped with an appreciation of the problems involved and a knowledge of the techniques which can help them. This type of research is therefore worthwhile both for the results it can produce, and also as a contribution to training men effectively for the chemical industry.

Further Reading

1. The strategy of process development, D. H. ALLEN and R. E. JOHNSTONE, *A. I. Ch. E.—I. Chem. E. Symposium Series 7* (1965) 67.
2. Two new tools for project evaluation, D. H. ALLEN, *Chem. Eng. 74 (14)* (1967) 75.
3. Credibility forecasts and their application to the economic assessment of R. and D. projects, D. H. ALLEN, *Operational Res. Quart. 19 (1)* (1968) 25.
4. A practical study of the accuracy of forecasts in novel projects, D. H. ALLEN and P. J. NOVEMBER, *Tripartite Chemical Engineering Conference, Session 12, Montreal (September 1968)*.
5. The concept of credibility and its use in forecasting for novel projects, D. H. ALLEN, *ibid.*
6. Optimum information accumulation patterns in the development of new chemicals; D. H. ALLEN and P. J. NOVEMBER, *Symposium on Venture Analysis, I. Chem. E. (November 1968)*.
7. Optimal selection of a research project portfolio under uncertainty, D. H. ALLEN and T. F. N. JOHNSON, *I. Chem. E. Symposium, Nottingham (April 1969)*.