Hitotsubashi Journal of Commerce and Management 26 (1991) pp. 29–42. © The Hitotsubashi Academy

SEQUENTIAL DECISION MAKING: SOME NON-TECHNICAL APPLICATIONS OF DYNAMIC PROGRAMMING*

J. M. NORMAN

This lecture is one of a series inaugurated in 1985, when the Jerwood programme was established. The programme is intended to contribute to the promotion of Anglo-Japanese cultural understanding through the exchange of lecturers with and through Hitotsubashi and Sheffield Universities, and the Jerwood Fellow at Hitotsubashi is normally asked to speak on some subject related to British culture, society, literature or history.

My own field is Operational Research, sometimes known as Management Science, the application of scientific method (very broadly interpreted) to management problems, particularly those arising in industry. You may well wonder what connection there can be between a technical subject like this and British culture, or any other country's culture for that matter. Consequently, this lecture is non-technical, and the examples used to illustrate points will be concerned, if not with British culture, at least with aspects of the British way of life, such as motoring, marriage, the pub, games and sports.

The word "sequential" is one to be emphasised so as to make the point that the decisions we make at any time affect the decisions we make, or are in a position to make, in the future. It is an obvious point that any decision is part of a chain, being both followed and preceded by other decisions. It is a commonplace which we all may acknowledge but whose application most of us often forget. It is very tempting to "take no thought for the morrow." One of the aims of this lecture is to illustrate how often and in how many places, the idea of a decision as one of a sequence of decisions may arise. Another aim is to suggest some ideas which may help in deciding what to do in situations of this kind.

The first illustration comes from personal experience, but it is one which every car driver must have had in mind at some time. You are driving along in your own car when you bump into another car. The accident is your own fault and you are fully insured. Should you claim? Of course, it depends. It depends on how likely you are to have further accidents. It depends how serious or costly these accidents are likely to be. It depends on the size, if any, of your no-claim bonus.

Questions of this kind have been studied in many papers, one of the most accessible being Hastings (1976). Hastings assumed that accidents occur randomly, that is in a Poisson process, at some rate which depends on the individual—how he drives and how much he uses the car. He assumed also that when an accident occurs, the subsequent repair cost

* This paper is, with minor changes, the script of a Jerwood lecture given at Hitotsubashi University on 19 July 1990. Most of the changes have been made to take account of the reduced number of illustrations possible in the printed version.

is a random variable which has a negative exponential distribution, an assumption which is in accord with what actually happens, at least for a certain class of motor accidents (Drinkwater and Hastings (1967)). For convenience, he took the maximum premium for a car owner to be £100 per annum and considered an actual policy which allowed net premiums to take values of 100%, 75%, 60%, 50% and 40% of this maximum. What happens to the premium from one year to the next depends on the number of claims made during the year, as shown in Table 1. The broad terms, you lose up to three years' bonuses if you make one claim and all your bonuses if you make two or more.

When you have just had an accident, let us assume that you can quickly find out what it will cost to put things right. You may now consider whether you should claim from the insurance company or pay for the repairs yourself. You know, of course, what your net premium was, last time you paid it, so you could now consult Table 2, which shows the optimal no-claim limits for a range of values of the accident rate. You should claim on the insurance company only if the repair cost is greater than the amount shown in the table.

There are a number of assumptions implicit in the derivation of Figure 2. One is that the average cost of a repair is equal to the basic premium. Another assumption is that the no-claim limit is the same irrespective of whether or not a claim has already been made in the year. These are grand, even heroic, assumptions, but the figures in the table are still of interest. It is noticeable that the no-claim limit is highest, for a given degree of accident-proneness, when you are halfway to maximum no-claim bonus. When you are paying the basic premium anyway, it does not matter too much if you claim on your insiurance—you have only delayed the transition to maximum no-claims bonus by a year. When you are on maximum no-claim bonus, then it is not too serious a matter (under this policy) to stay above this level for two years (or more than two years if you are relatively unlucky).

TABLE 1. NETT PREMIUM LEVEL SUBSEQUENT TO A CLAIM (ROYAL INSURANCE POLICY)

Current nett premium £	Subsequent nett premium: Zero claims £	One claim £	Two or more claims £
100	75	100	100
75	60	100	100
60	50	100	100
50	40	75	100
40	40	60	100

TABLE 2. OPTIMAL NO-CLAIM LIMITS FOR A RANGE OF VALUES OF THE ACCIDENT RATE

Current nett premium £	Accident rate 0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8
100	60	50	50	50	40	40	30	20
75	100	100	90	80	80	70	60	60
60	120	120	120	120	110	100	100	90
50	70	80	80	80	90	90	80	80
40	40	40	50	50	60	60	60	60

The foregoing decision problem may not only be of some practical interest but can be used for a methodological purpose, that is, to illustrate the points that need to be established before the calculations can be carried out and the optimal claims policy established. In the problem we can identify four aspects:

1. the situation you are in:
   the premium you are currently paying
   the cost of putting right the damage which has occurred
   (in technical terms we might call this the state of the system)
2. the alternative courses of action open to you:
   to claim on the insurance company
   or to pay for the damage yourself
   (in other words the decision you can take)
3. what each of the alternatives will cost you in the short term:
   either nothing (if you claim)
   or the repair costs (if you don't)
   (these are the immediate or short-term costs)
4. the way each alternative transforms the state of the system:
   to an increased premium next year (if you claim)
   or to the same or a reduced premium (if you don't)
   (technically, we call this the system transformation)

There is one further aspect of the problem which has not been mentioned: the objective. In fact, the objective in the formulation discussed was taken as being to maximise long-term average costs. This is not the only possible sensible objective: if, in the short term, you can't afford to pay the repair costs, then you *have* to claim on your insurance, even though you might expect to do better, in the long run, by not claiming. One theme of this lecture is to show how identifying the characteristics of a problem and identifying the objective can help our judgement in problems where some decision has to be made.

The next illustration is concerned with a different kind of objective—not a process in which we try to minimise the sum of period-by-period costs or to maximise the sum of period-by-period profits, but one in which we try to maximise the probability of achieving some desired state—in this case, matrimony, to the best spouse possible. The simple (if biased) model considered is that throughout his adult life, or at least in its early stages, a man meets a succession of girls. It is not unreasonable to suppose that he meets them one by one, that he forms some sort of impression of each one he meets and considers seriously as a potential partner in marriage before passing on to the next. It is less reasonable, but not out of the question, to suppose that he may propose marriage to any girl in the sequence, and that a proposal, once made, will be accepted, and a proposal not made will be counted by the girl as a rejection. There can be no going back.

This is clearly a sequential decision process: a succession of girls appear in our life, to each of whom we must either propose or not and the transition is either from the single to the married state or to remain single. The costs and returns of this process are rather more difficult to establish. The objective is clear cut: to propose to the best girl—of all those whom we have already met and of all those whom we shall meet later if we con't propose to this one. The best or nothing. All we need to be able to do is (subjectively) to choose the best out of any group.

It is not very difficult to show that the optimal policy is to study the first 37% of all the girls whom you expect to meet but to take no action, that is, not to propose marriage to any of them. Afterwards, you should propose to that girl whom you then meet who in your opinion is the best of all you have met so far, including those of the initial 37%. Clearly this policy cannot guarantee you success, in the sense of guaranteeing that you will be able to propose to the girl you like the best—if you follow it, you may find yourself not proposing to anyone—but it does maximise the probability of your being able to do so.

It is tempting to draw conclusions of a sociological nature from this result. One argument is that the sequence of marriageable girls begins when a man is aged (say) 16 and more or less peters out when he is aged (say) 40. On the assumption that the rate at which the girls appear does not change over the years, then a man should wait until he is 25 (that is until 37% of the years between 16 and 40 have elapsed) before considering proposing marriage. It is a common observation that most Englishmen do not follow this policy and consequently it might be argued that most Englishmen marry too young. It is noteworthy that currently, the average age at first marriage of Japanese men is 28. Nevertheless, the model, useful though it is, cannot be used to support this sociological conclusion. Most men do not experience a constant stream of potential wives from 16 to 40. More seriously, there are difficult questions relating to (a) the validity of the objective, (b) the possibility in practice of choosing the girl you think best and (c) the possibility of changing values and adaptability. That the model discussed is not necessarily applicable to every adult male who has ever lived is easily shown. King Henry VIII of England had six wives, one after the other. Somewhere in Africa is a man who has more than two hundred.

In decision making, a model is necessary to justify a chosen action. Whether a particular model is appropriate is a matter of judgement. To illustrate this point, consider the following problem which was discussed in the New Statesman (Young 1978). The date is some years ago, when public houses were compelled, by law, to close at 2 pm for the rest of the afternoon. It is lunchtime and you are in your car, with a lady companion, in an unfamiliar part of the country, and you are looking for a pub. At each pub you come to, you have to decide whether to stop and go in, or to drive on in the hope of finding a better one later. The cost of driving on is not just the variable cost of petrol and wear and tear on the car but also the distress to driver and passenger as valuable drinking time slips away. The problem characteristics are well marked. We can recognise the decisions to be taken, the costs, and the way the system moves from one state to another over time. What is not well marked is the objective: clearly we want to get a drink before closing time, but how much before closing time? How important is drinking time to us? In fact the author of the article took the objective to be to maximise the probability of having a drink in the best pub, and he suggested a near-optimal policy, in which precision was sacrificed to the convenience of approximation.

The policy suggested was not to go into the first pub but to go into the best pub after that. It is true that this *is* a policy which is not far from optimal, with the onjective stated, and on the assumption—comparable to the one used in the previous example where we considered choosing a wife—that we know nothing about pubs before we start. This assumption might well apply, say, to a couple from a far off country (Japan, perhaps) who have heard something about the English pub, but, being new to England, have never been inside one before. In such a case, a possible conversation on seeing the first pub might be: "So that's a pub, well, well. Let's have a look at one or two more to see what some others are like." Such a remark would be congruent with the pursuit of optimality: we need that 37% to see what the general run of pubs is like. But a remark of this kind could hardly be made by an English couple, who must already have a very good idea of what pubs are like. The point is that, in the latter case, we can do more than merely put the pubs in an order of merit (in our eyes) as we see them—we can grade them. In this instance we already know quite a lot about pubs, and this ought to be taken into account when we consider whether or not to enter a particular pub.

It is not a straightforward matter to evaluate a pub. To determine what people *in general* look for in a pub is a matter for marketing rather than operational research, and has been the object of much study (Meidan (1978)). The consideration of an individual's pub selection criteria effectively involves us in the construction and calibration of multi-attribute utility functions. Devotees of real ale have frequently to put up with uncomfortable surroundings: for them, the quality of the beer may more than compensate. The judgements of an individual on each of beer quality, pub decor, expected degree of crowdedness and so on are compounded (almost always unconsciously) in a measure which expresses the utility to him, at that time, of a particular pub. Some people argue that this kind of composite evaluation is impossible, but we do it (implicitly) whenever we enter, or decline to enter, a pub.

We have identified the four characteristics of the decision situation:

the state (where we are)
the decision (where do we go from here?)
the costs or profits (what will it cost to go there?)
the transitions (where will we get to next?)

Over and above the characteristics of the problem which we face at each stage, that is, at each point in time when a decision has to be taken, we have our objective, the criterion which enables us to judge the superiority of one decision over another. The four problem characteristics, together with the objective (or the "criterion of optimality," as it is sometimes called) together form a helpful way of looking at sequential decision problems. Indeed it would be difficult to get very far with the analysis of such problems without first of all clearly identifying these aspects. There are two methods of analysing such problems which are widely used today, which have been developed over the last thirty or so years, principally in the United States, but to some extent in the United Kingdom, Japan and elsewhere. These two methods are dynamic programming and decision analysis.

Dynamic programming is associated with the work of Richard Bellman, who began and developed his ideas when working with the RAND Corporation in the 50's and 60's. A prolific and inventive writer, the book called simply "Dynamic Programming" (Bellman (1957)) is still a fruitful source of problems and research ideas. In this book Bellman first stated his "Principle of Optimality," which, embedded in a functional equation appropriate to each problem, is the basis of dynamic programming. Bellman stated the principle thus: "An optimal policy has the property that, whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision." Once you understand what is meant by this unwieldy sentence, you wonder how a statement so obvious can be so fruitful in its applications.

It is difficult to state the principle clearly. Professor White has stated it more concisely: "All contractions of an optimal policy are themselves optimal" (White 1969), but you have to know what contracted policies are to understand this, and it is not a trivial matter to define them. A definition in an engineering text runs as follows: "If you don't do the best you can with what you've got, you'll never know if you'd have done as well as you might with what you should have had."

Although it is hard to express the principle in simple terms, its application in practice is often obvious. After damaging our car in an accident, we should not be concerned with what has happened in the past, except to use it as an indication of what might happen in the future, that is, as regards the possibility of having further accidents. Our concern should be with the situation we are in now, and our intention should be to choose wisely now and in the future. That (say) we claimed on the company two years ago is irrelevant to the present decision: it has already been taken account of in the premium we paid last time.

In an example like this, the good sense of the principle of optimality is so obvious that we may suspect that its importance is trivial. On the contrary, its importance is profound: it is remarkable than an idea so seemingly obvious can generate a technique as computationally powerful as dynamic prohramming.

That the principle of optimality is straightforward does not imply that its application in dynamic programming is straightforward too. Even in the car insurance claim problem, the solution given earlier was obtained using a number of simplifying assumptions. To make the example more realistic needed a considerable extension of the computational effort. In further work carried out on the problem (Norman and Shearn (1980)) allowance was made for the possibility of more than one accident occurring in the year between successive insurance premium payments, and the unknown average cost of damage was set at more than one value. Account was also taken of the time in the premium year at which the accident occurred. The upshot of all this labour was that Hastings' rules, when re-interpreted in terms of the premium which will be paid next year if no more claims are made, turned out to be very close to optimal. More than that, the commonly used rule of thumb, which suggests claiming only if the costs associated with the accident are greater than the total of the increases in premiums which will have to be paid if a claim is made, turned out to be a policy nearly as good as the optimal policy, with an expected cost generally less than 2% greater than the expected cost of the optimal policy.

The problem of how best to balance accuracy or realism against the computational load is ever present in operational research, but it is particularly acute in dynamic programming. Were it not for the existence of digital computers then dynamic programming could never have reached the position it now has as a powerful problem-solving technique: the computational burden in all but the very simplest of sequential decision problems is such that a computer is needed to cope with it. In some problems the burden is such that adjustments have to be made to the model constructed to describe and define the problem in order to reduce it to a manageable size. This should not be taken to imply that simple problems are not necessarily worth solving. Simple solutions are often worthwhile and they are more likely to be implemented than complicated ones.

The other method of analysing sequential decision problems is usually called decision analysis. In the United States, its development was associated with Professors Raiffa and Schlaiffer of the Harvard Business School and in the United Kingdom with Professor Lindley of University College London and Professors Moore and Thomas of the London Business

School. I believe that these authorities have had a much greater influence on practical decision makers, particularly businessmen, than the authorities on dynamic programming. In a way, this is curious, in that in many cases, the two approaches may be regarded as two aspects of the same method. For a particular problem, one may be more convenient, or more apt, than the other, but the implied methodology of both is the same.

The application of decision analysis often takes a schematic form and is easy to illustrate. A well known example (Fig. 3) shows a very simple form of the decision faced by

Fig. 1

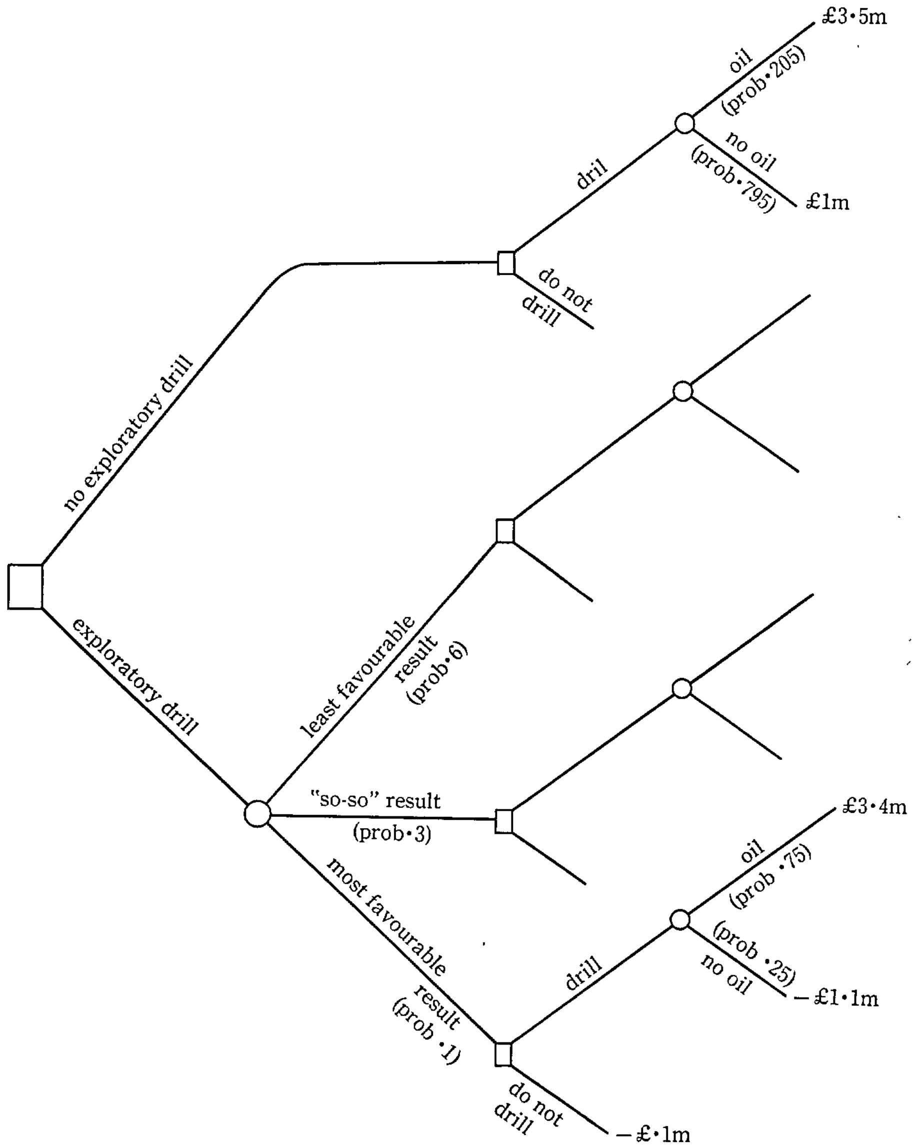

an operator who is drilling for oil in the North Sea. Actions and events are shown by lines, decision lines emanating from squares (decision nodes) and event lines from circles (event nodes). The operator has the problem of deciding whether or not to seek further information about the possibility of finding oil by carrying out an exploratory drilling. The argument in favour of doing so is that the results will make more precise the probability of his finding oil without going to the expense of a full scale drilling operation which might turn out to be not worthwhile. Obviously, a number of assumptions are made. It is taken for granted that the operator can assess all the probabilities, that he can estimate all the costs and returns, and that there are no difficulties in ranking the possible consequences. Given all the data appropriate to the decision tree, it would not be difficult to work back from the "twigs" of the tree, along the "branches," to the "trunk". If the operator's objective is to maximise his expected return, then it may be that he should carry out the exploratory drilling, and drill full scale unless he gets the least favourable result. We could derive this policy using a more abstract dynamic programming formulation but the underlying ideas would be the same. Notice, for example, that if the operator has made an exploratory drilling and has observed the most favourable result, then, at this point, the decision he takes should not be influenced by his financial outlay on the exploratory work, but only by the revised probability of finding oil. The principle of optimality has as central a place in decision analysis as it has in dynamic programming.

I have used this well-known example advisedly. That businessmen, at least in the United Kingdom, are now more receptive to the idea that information about uncertain events can be quantified is in large measure due to their becoming accustomed to probability ideas through reading and talking about North Sea oil. North Sea oil may not have revitalised the British economy as much as was hoped, but it has, at least, been a help to academics in decision analysis, helping them to convince businessmen that uncertainty can be quantified.

There has not, I think. been the same difficulty in the United States, where businessmen are hungry for numbers: not just data on costs and quantities but on probabilities also. Most operational research workers visiting America from Britain appreciate their change of environment when they hear an American weather forecaster. Used to variations of "Bright periods with scattered showers, often of rain, sleet or snow, which may be heavy and prolonged in places" they hear instead "The probability of precipitation is . . . " followed by a number quoted to within 10%. This kind of forecast is not uncommon in Japan. It is a useful exercise for undergraduates and others to consider how the accuracy of such forecasts can be checked. It is a more difficult exercise to determine what such statements imply about the ability of the weather forecaster.

Decision analysis, in the restricted sense of analysing decisions and events graphically, is not a new technique and its use is certainly not confined to business. Examples of its use in the analysis of chessboard positions are common (Kotov (1971), for example). Chess has been described as a trivial game. Most people would agree that noughts and crosses is trivial: nobody need lose a game of noughts and crosses if he knows the rules and plays sensibly. I have been told that in the variant of the game of Go in which each player tries to make a line of five counters, the first player should always win if he plays optimally. A remark of the same kind must apply to chess, but nobody has yet determined what the remark is. Can White necessarily force a win? Should Black always be able to draw? There is a sense in which there is no uncertainty and no risk in chess; this does not imply, of course, that there is no luck, at least in a game played between human beings.

Games of skill, such as snooker and football, or athletic competitions, such as throwing the discus or weight lifting, are a different matter. These games are certainly not trivial and much effort, though not as much as in chess, has gone into attempts to analyse them. An example of such an analysis (Ladany et al. (1975)) relates to the long jump. In this event an athlete usually has three attempts and his best jump is the one that counts. For a jump to be recognised it must begin with both feet behind the official take-off line, although the distance counted is measured from the take-off line itself. An analysis of the jumping style of a particular United States athlete showed that the length of his actual jump followed a Normal distribution with known mean and variance, and that the accuracy of his take-off —the distance between his actual take-off point and the official take-off line—also followed a Normal distribution with known mean and variance. The analysis showed that if this athlete wished to maximise his expected distance counted he should aim to take off for his first jump from an imaginary line about 1 1/2 inches before the official take-off line, for his second jump (if the first jump was not valid) from a line 5 1/2 inches before the official take-pff line, and for his third jump (if the first two jumps were not counted) from a line 10 inches before the official take-off line.

A long-jumper's run-up is usually of the order of thirty yards. To suggest that an athlete, approaching his top speed, can place his feet to an accuracy of two-and-a-half inches is stretching credulity. Overemphasis on this kind of recommended policy can bring operational research into disrepute. The example furnishes another exercise in determining an appropriate objective. A long-jumper rarely, perhaps never, wants to maximise his expected distance jumped. He is more likely to try to maximise the probability that his best scored jump is greater than a specified distance, as indeed was suggested in a further analysis (Ladany and Sphicas (1976)). If this is the case, then his optimal policy is to aim for a line at a constant distance from the official take-off line and this strategy can be operated easily enough through practice in the athlete's training.

A similar analysis would hold good for the field events in which the athlete is limited to three attempts, such as the javelin and the discus. Rather more complex are those events in which the athlete can choose the precise feat he will attempt. For example, in the high jump or the pole vault a competitor can specify the height of the bar. If he sets it low he runs a small but not negligible risk of not clearing it in three jumps, but gives himself practice in readiness for greater heights later on in the event. Added to this, he may be influenced psychologically by the tactics of his opponents in that they too can decide where the bar should be set. The most extensively analysed event in this category may be weightlifting, where the athlete becomes more tired by attempts to lift weights and hence is restricted in the total number of attempts to lift weights he might anticipate being able to make (Lilien 1976)).

Perhaps the most prolific source of problems is orienteering, where the objective is generally to find one's way to a succession of control points in the shortest time. This is a problem to which it has been possible to provide satisfactory solutions surprisingly simply. In an analysis of one leg of the 1974 Vaux Lake District mountain trial (Hayes and Norman (1984)) the ordnance survey map of the area was divided in a 250 metre grid. At each grid point, the spot height and the nature of the terrain (rough or especially rough) were noted.

Using commonsense rules of thumb for running uphill, downhill, and on the level, the time taken to run from any grid point to an adjacent grid point could be determined easıly, and thus the optimal route established, using perhaps the simplest of all functional equations in dynamic programming.

The "optimal" route turned to be significantly different from the route taken by the winner (Fig. 2). The reason is that the winning athlete needed a drink of water at that stage in the race and took a longer route in order to pass by a spring which he knew existed but which is not shown on the map. The optimal route was the one chosen in the race by my colleague. Optimal routes from various other starting points, together with hte standard times required to run along them to the terminal point, were also determined (Fig. 3). Two points are worth noting. First, from some neighbouring points, the optimal routes are very different. Consequently, a starting point in such a neighbourhood would make a demanding test of the route-finding abilities of an orienteer. Second, the points occurring where several routes merge should be suitable places to site services such as rescue

FIG. 2

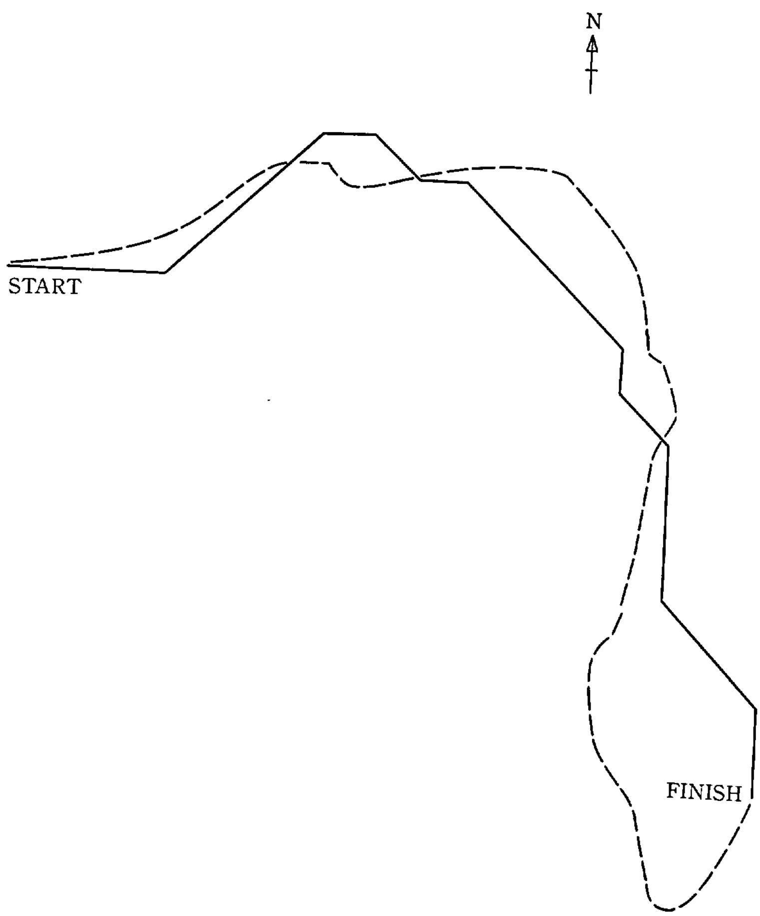

FIG. 3

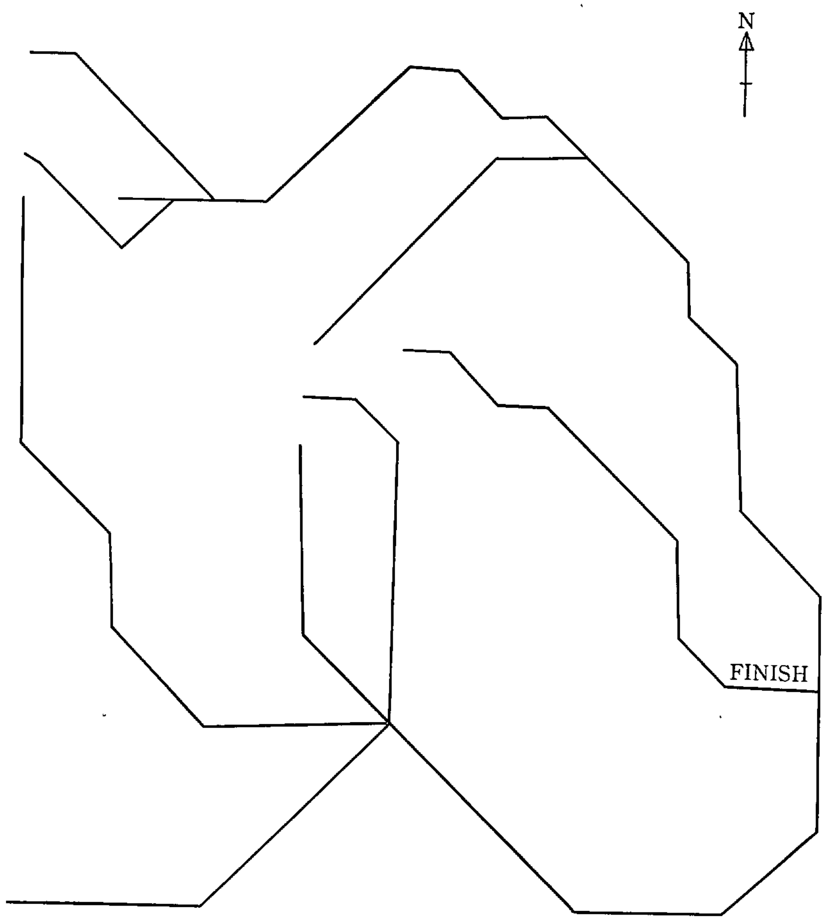

teams. This is an example of the operational research approach helping to solve other problems as well as the one originally posed.

This work has encouraged us to investigate how comprehensive the state description needs to be. We have looked at the performance of athletes on the Bob Graham round, who cover within 24 hours a 72 mile circuit in which they climb 27,000 feet. The distance is equivalent to three marathons, and the height climbed is equivalent to running up Mount Fuji from the bottom to the top 2 1/2 times. The main factors seem to be simply position and time of day—you can't run so fast at night when it is more difficult to see where you are going. How long the athlete has already been running does not seem to be an assential variable in the state description. This somewhat surprising conclusion is in accord with what many athletes have said about the importance of even-paced running in long distance events, of which marathons are a good example (Phillips 1983).

Some orienteering events are much more complex. Usually, the competitor has a map, but in certain events, for experienced orienteers only, the competitor has to find the next control point, where a map will be fixed, and this he will have to memorise in order to find the control point after that. Another complex event is one in which the competitor is handed at the start a map on which a number of controls are marked, each having associated with it a score. The competitor's task is to navigate himself to such control points as he thinks appropriate in order to maximise his total score within a set time period, usually one hour.

All of these problems are sequential decision problems and in principle, all of them are capable of solution by dynamic programming methods. That many of them have not been solved is due partly to there being no financial incentive to look for a solution and partly to the complexity of the individual problem. It is worthwhile looking at these problems, however, because the methods of solution appropriate to them can also be employed in problems which occur in the business world. As well as this, it is much easier to test approximate, or near-optimal solutions, in an aiea of application such as sports tactics than it would be in industry. No athlete is going to mind very much if he does less well than he might in a training event but a chief xeecutive might be understandably irritated if his firm loses money because of experiments carried out at the suggestion of his operational research man who wants only to get more data.

In this lecture I have tried to illustrate the variety of sequential decision problems and thus the scope for application of the ideas of dynamic programming and decision analysis. However, these ideas can be difficult to apply and my two final remarks will be concerned with two ways in which it may be possible to sumplify sequential decision problems and which thus may make them easier to solve.

The first remark is that sometimes what one should do—the optimal policy—is sometimes very simple. Even complex problems often have easy solutions. At one time, the Central Electricity Generating Board were concerned to find the best way of developing electricity substations so as to meet the demand for electricity in the areas they served at minimum discounted total cost (White and Norman (1969)). To have a bigger substation thean necessary meant large costs due to excess capacity (Fig. 4). To have that size of

FIG. 4

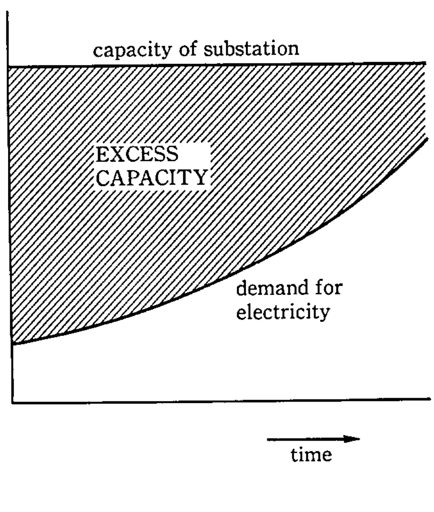

substation which would just meet the demand at any time would reduce excess capacity but would result in expense through frequent changes in the capacity of the substation (Fig. 5). What was required was clearly a balance between excess capacity costs and expensive capacity adjustment. The calculations of the relevant costs were complicated but it was important to get them right since the figures involved were in millions of pounds. It became apparent, after not many minutes of computer time, than an optimal substation sequence always had one characteristic. namely that it was never worthwhile to increase the capacity of a substation until it could no longer meet the demand. The recognition of this characteristic considerably reduced the computational effort; it was also a very useful rule of thumb for the decision makers of the Generating Board.

The second remark is that very often you don't need to look very far ahead. The investment manager of a certain insurance company was concerned with the optimal use of short term funds: whether to hold them as cash or to invest in 2-day or 7-day loans (that is, loans which could be recalled on 2 days or 7 days notice). The longer the period of notice, the greater the interest, but the greater the risk of incurring overdraft charges through the insurance company being unable to meet its very short term obligations immediately from its current account. It turned out (White and Norman (1965)) that the investment manager was following a policy very close to optimal, and perhaps surprisingly, the form of the optimal policy turned out to be the equivalent to a one-period ahead policy. That is, the manager effectively considered what was going to happen in the next week, and that week only. What was going to happen after that, unless it was going to be such a radical change that the dynamic programming model would no longer apply, could be ignored. It was as if he were putting into practice the words of the hymn "I do not ask to see the distant scene, one step enough for me."

Neither of these remarks will surprise anyone who plays pool or snooker. In snooker, it is good play and simple policy to keep the red balls close together but not tightly compact and to try to make your breaks at the top of the table. Second, it is a good idea to look just one shot ahead and to know where your cue ball will finish at the end of your shot. It

FIG. 5

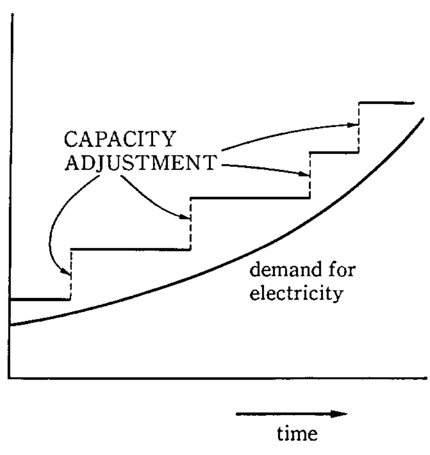

is a pity that although dynamic programming is such a useful method of analysis and can help you to discover what you ought to do, it cannot of itself give you the ability to put that policy into practice.

UNIVERSITY OF SHEFFIELD

REFERENCES

Bellman R E (1957) Dynamic Programming. Princeton University Press, Princeton, N J.

Drinkwater R W and N A J Hastings (1967) An economic replacement model. Opl Res Q, 18, 121–138

Hastings N. A J (1976) Optimal claiming on vehicle insurance Opl Res Q, 27, 805–813

Hayes M and J M Norman (1984) Dynamic programming in orienteering: route choice and the siting of controls. J Opl Res Soc 35, 9, 791–796

Kotov A (1971) Think like a grandmaster. Penguin, Harmondsworth

Ladany S R, J W Humes and G P Sphicas (1975) The optimal aiming line. Opl Res Q 26, 31, 495–506

Ladany S P and G P Sphicas (1976) Dynamic policies in long Jump. Management Science in Sport 4, 113–124

Lilien G. L (1976) Optimal weightlifting. Management Science in Sport 4, 101–112

Meidan A (1978) Pub selection criteria. Proc Marketing Education Group Conference, United Kingdom

Norman J M and D C S Shearn (1980) Optimal claiming on vehicle insurance revisited. J Opl Res Soc 31, 2, 181–186

Phillips B (1983) Is the two hour marathon a possibility? Running Review 12, 24. Ron Hill Sports, Hyde, Cheshire

White D J (1969) Dynamic Programming. Allen and Unwin, London

White D J and J M Norman (1965) Control of cash reserves. Opl Res Q 16, 3, 309–328

White D J and J M Norman (1969) An example of problem embedding in deterministic dynamic programming. Opl Res Q 20, 4, 469–476

Young P (1978) The Young system: decision making without tears. New Statesman and Nation, 28 Nov.