A RAND REPORT

PREPARED FOR

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION

WEATHER INFORMATION AND ECONOMIC DECISIONS
A PRELIMINARY REPORT

Richard R. Nelson
Sidney G. Winter, Jr.

RM-2620-NASA

August 1, 1960

This research is sponsored by the National Aeronautics and Space Administration under Contract No. NASw-103, monitored by NASA's Committee on Long-Range Studies. This report does not necessarily represent the views of the National Aeronautics and Space Administration.

The RAND Corporation
1700 MAIN ST. • SANTA MONICA • [CAL]IFORNIA

SUMMARY

Most improvements in knowledge come only at a cost, and better knowledge about the weather is no exception to this generalization. What are the chances that the effort now being undertaken to develop a meteorological satellite will turn out to be worthwhile? Our offhand answer to this question would be "very good." But the reasons we might give for our answer would probably not be very convincing to someone who did not have a good deal of faith in technological progress. To give an answer that is more than just an opinion, it is necessary to answer these two questions: (1) To what extent will satellites result in a better capability to make weather predictions? and (2) even assuming that much better predictions will be possible, what will this improved knowledge be worth? We cannot answer the first question, because we do not know for certain what particular types of data meteorological satellites will be able to collect, nor to what extent the additional information will improve our ability to forecast the weather. But if we postulate the improvement in predictions that could take place, we can answer the second question, and this we attempt to do. Although our analysis will not yield any exact figure for the economic benefits from satellite observations, it should provide some quantitative indication of what their value might be.

Many individuals and business firms must make a wide variety of decisions on the basis of weather forecasts. An incorrect forecast may mean heavy losses, for expensive protective measures can be taken needlessly or damage can occur which could have been avoided. Better forecasts mean fewer losses. Our analysis deals with the question of how much

better off individuals or business firms would be with improved forecasts. In particular, we shall try to calculate dollar values.

RAND's study of the economic benefits from improvements in weather forecasts falls into three major parts. First, we are developing a general analysis to determine the value of a forecasting service to a single economic decision maker. This analysis shows how the value of the forecasting service depends on the economic characteristics of the decision problem and on the quality of the forecasts, and consequently permits us to determine the value of any particular improvement in weather forecasting to a particular decision maker. Second, we are modifying this general analysis in order to provide more convenient tools for dealing with certain classes of problems in which decisions are made sequentially on the basis of weather information. Third, we are examining (a) the problem of how the weather forecaster can exploit his understanding of weather phenomena in order to make forecasts of the greatest possible value, and (b) the criteria which are relevant in solving this problem.

If we specify what information the decision maker would have in the absence of forecasting, it is possible to calculate the value of a forecasting service and to compare one such service with another. Our analysis is based on the prior work of J. C. Thompson, G. L. Brier, I. Gringorten, and other meteorologists, and on the general developments in the field of statistical decision theory and the economics of information. This general analysis is useful, not only for providing a framework within which the value of particular forecasting services or their improvements can be determined, but also for the general insight it affords into the determinants of the value of information about the weather. It also provides a

means for valuing perfect weather forecasts, and therefore puts an upper bound on the economic benefits to be derived from improved knowledge of weather phenomena. Of course, in many situations it is difficult or impossible to assign a dollar value to an improvement in weather forecasts, and then, too, there are the intangible benefits that one encounters in any increase in the knowledge of natural phenomena. Our analysis, therefore, does not provide a framework for examining the full range of benefits which might be expected to flow from the meteorological satellite program. But it does provide a framework for investigating some of the important benefits.

Though we have not gone as far as we had hoped to in making empirical studies, we have applied the methods that we have developed to three different kinds of firms. These particular case studies were chosen for their analytic convenience rather than for their social significance; nonetheless, the empirical results suggest that the value of improvements in forecasts can be considerable. For example, in studying the value of present weather forecasts to a trucking firm and the value of improved forecasts, we found that present forecasts permitted costs associated with inclement weather to be reduced by about one-third, and that with better forecasts, further reductions by as much as 80 per cent of the remaining cost might be possible. Very similar results were obtained from a study of a motion-picture firm. On the other hand, our third study, which involved a roofing-construction firm, did not indicate commensurate benefits.

While the results of three studies can hardly be regarded as conclusive, we believe that further studies will show that improvements in weather forecasts can be important even in activities that would not ordinarily be regarded as extremely sensitive to the weather.

Though the approach taken in these studies is adequate for a very wide range of decision problems arising in activities affected by the weather, there are other types of decision problems -- where decisions are made sequentially on the basis of a continuing flow of information about the weather -- which become computationally intractable when placed within the general framework. The second part of our work is aimed at developing special purpose models which exploit the specific structure of particular problems in order to avoid this computational difficulty. In this connection, we have analyzed a set of problems in which a decision maker must choose an opportune time for attempting a particular job. The decision to attempt the job in what turns out to be bad weather carries with it some loss, but there are costs associated with delay in getting the job done, either in the form of a charge per day, or a penalty after a certain deadline passes, or both. Good forecasts enable the decision maker both to avoid the losses of doing the job under unfavorable conditions and to exploit favorable conditions when they occur. This framework has been applied to a roofing-construction problem.

In the third major portion of our study, we consider the problem of how the forecasts themselves should be made if maximum economic benefit is to be derived from the knowledge of weather phenomena which underlie the forecasts. Consider, for example, the problem of deciding under what conditions a forecast of rain should be made. The forecaster encounters, over a period of time, a wide variety of conditions in which he can make a forecast of rain with varying degrees of confidence. By forecasting rain only when the signs of rain are relatively unequivocal, he can make a rain forecast which will actually be followed by rain a very high percentage of

the time -- but at the price of having rain occur frequently when he forecasts fair weather. We show how the needs of consumers of the forecast should influence the forecaster in making his decision. An argument for making forecasts in probability form is presented with the qualifications to that argument. We use J. C. Thompson's objective scheme for forecasting rainfall in the Los Angeles area, as a basis for some empirical computations.

Thus our work to date can be described as developing appropriate analytical frameworks and testing their practicality against easily obtainable data. In later studies we shall apply these methods of analysis to cases selected both because of their major economic value and because of the likelihood that the relevant forecasts may be improved as a result of data obtained from meteorologically instrumented satellites. The problem of hurricane warnings seems to meet both these requirements and will be the subject of our next study.

CONTENTS

SUMMARY iii

I. THE GENERAL FRAMEWORK 1

INTRODUCTION 1

AN EXAMPLE 4

TOWARD A MORE GENERAL ANALYSIS 21

II. SELECTED CASE STUDIES 25

INTRODUCTION 25

THE VALUE OF METEOROLOGICAL INFORMATION TO THE MOTION PICTURE INDUSTRY 28

THE VALUE OF METEOROLOGICAL INFORMATION TO THE CONCRETE LAYING BUSINESS 39

III. WEATHER FORECASTS AND ACTIONS WHICH CAN BE DELAYED 43

INTRODUCTION 43

A LINEAR POSTPONEMENT MODEL 44

AN EMPIRICAL EXAMPLE: ROOFING CONTRACTING 50

A POSTPONEMENT MODEL WITH A DEADLINE 53

IV. THE MAKING OF FORECASTS: OPTIMAL FORECAST STRUCTURES 67

INTRODUCTION 67

A DIGRESSION ON THE CONCEPT OF FORECASTING 68

CRITERIA FOR OPTIMAL FORECAST STRUCTURES 73

TWO EXAMPLES OF THE DETERMINATION OF OPTIMAL FORECAST STRUCTURES 77

MULTIPLE ACTIONS, MULTIPLE USERS, AND PROBABILITY FORECASTS . 97

V. THE DIRECTION OF FUTURE RESEARCH 105

APPENDIX 1 TO SECTION I 109

APPENDIX 2 TO SECTION I 113

APPENDIX TO SECTION IV 119

I. THE GENERAL FRAMEWORK

INTRODUCTION

If it were true that although everyone talks about the weather no one does anything about it, the science of meteorology and the art of weather forecasting would be hardly more than intellectual exercises. If the only sensible response to advance warning of adverse weather were to grin and bear it, then meteorological information would have no value. In some situations fatalism may be justified. For a farmer who has just planted a tender crop, information that sub-freezing temperatures will occur during the next few days may be of no more value than bad news in general. But in many situations people can do something about the weather, and meteorological information has value because it helps them to do the right thing. For a farmer who has not yet planted his crops, a reliable weather forecast may avert disaster.

A wide range of human activities are sensitive to the weather:[1] industrial, agricultural, military, domestic, recreation, to name only a few. In the vast majority of these activities, men can do more than grin and bear it; the best action they can take differs as a function of the weather. For farmers in making their crop choices, in deciding when to plant and harvest, and whether or not to protect their crops against the possibility of adverse weather; for airlines in deciding their flight routes and

---

[1] For excellent studies of the effect of weather and climate on agricultural and industrial activities, see Climate and Man, 1941 Yearbook of Agriculture, U.S. Government Printing Office, 1941, and Industrial Operations Under Extremes of Weather, Meteorological Monograph No. 9, Vol. 2, American Meteorological Society, May 1957.

passenger loads; for public utilities in ordering fuel and in deciding whether or not to order standby capacity; and for transport companies in deciding when, how, and by what routes to ship, meteorological information has value. It has value because it helps people to make good decisions.

It must be stressed, however, that hurricane warnings have value not because hurricanes are destructive, but because with suitable warning people can take action to reduce destruction. Forecasts for sunny weather have value not because people like sunny weather, but because good forecasts permit people to plan ahead to take advantage of sunny days.

What do we mean by meteorological information? For the present we shall focus on the decision maker, and by meteorological information we shall mean information on weather or climate that affects their choice of action.[1] The type of information which has relevance varies from case to case. For some decisions the quantity of rainfall is relevant; for other decisions, snowfall, or temperatures, or cloudiness, or some combination of a number of meteorological elements. For some decisions the weather at a particular place and time is important; for others, the average weather over a relatively long period of time and over a wide range of territory.

The meteorological information that the decision maker has is difficult to quantify; nor is there, in general, a unique source of the information. Some information stems from the decision maker's own observations and is quite informal. Some of the decision maker's information is prepared and relayed to him by meteorologists. Meteorologists sometimes find it convenient to classify meteorological information into records of past weather,

---

[1] We shall expand and elaborate this definition later in this section.

observations of present weather, and predictions of future weather. Although this classification is less helpful when information is looked at from the point of view of the decision maker rather than that of the meteorologist, there certainly are many cases when the decision maker wants to know what the weather presently is, or what it actually was at a certain time and place. However most decisions must be made on the basis of predictions, and in this report we shall focus on the value of weather forecasts. When we consider information from the point of view of the decision maker, we will mean information about what the weather will be.

Climatology is a well-developed sub-science of meteorology which deals with past weather records in various systematic ways. Looking at climatological information from the point of view of the decision maker, we will mean the relative frequencies of different weather states at different times of the year. We shall assume that in the absence of other information the decision maker interprets these relative frequencies as probabilities of what the weather will be.

In this report we shall use the term forecast to denote predictions about what the weather will be which are made on the basis of observations and theory including more than climatology. Both climatological data and observations enter into the making of forecasts, but for the present we shall not be concerned with how the forecasts are made. No specific type of forecast or technique of making forecasts is implied. However, we assume that the decision maker can treat any specific forecast as a signal indicating a certain probability distribution of weather states.

As in many economic problems, absolute or total value is difficult to define and measure, and is irrelevant for most purposes. It is

usually more fruitful to examine marginal value. Given a particular decision maker and the kind of problem he faces, how much better can he do receiving a given kind of forecast than he could do if he had less accurate information? How much will improvements in forecasts be worth to him? This study will develop and apply methods designed to answer the above questions.

It must be stressed that we are looking at information from the point of view of the decision maker, not the meteorologist. For the meteorologist, information is the main material for analyses and forecasting. It is present and past weather, observations and records. It consists of weather at remote locations and of instrumental observations in which decision makers are not directly interested. It includes theory and various synoptic frameworks. The forecasts which are "information" for the decision maker are one end product of the meteorologist's art.

AN EXAMPLE

Before proceeding to a general analysis, it seems worthwhile to consider a relatively simple decision problem. There are a vast number of situations in agriculture and industry in which the range of possible actions the decision maker can take is quite limited, but it is possible for him to take protective action against adverse weather conditions at a cost. Thus a newspaper distributor, who has a standard routine for distribution, can wrap his papers in wax paper to protect them from rain. A storekeeper can tape his windows to protect them from a threatening hurricane. A citrus grower can light smudge pots to protect his fruit from frost. Many decisions are not of this protect-don't protect sort, and the analysis we will develop is capable of handling more general decision

problems. However, the simple protection problem is interesting and important, and serves as a convenient introduction to the more general analysis. It has been treated before in the literature on applied meteorology,[1] so the formulation may be familiar to many readers. In addition, the factors determining the value of meteorological information take on a simple and understandable form in this case.

More specifically, consider the problem of a dispatcher of a fleet of ten trucks. The schedule requires that all loading be accomplished on the day or evening before dispatch so that the trucks can leave early in the morning. All of the trucks are "flat racks" (uncovered trucks), and the merchandise carried is seldom of a sort which will be damaged significantly by light rain. It is customary to ship such merchandise -- building material, canned goods, etc. -- unprotected or protected only by normal packaging, if the chances of rain are judged to be slim.

However, a moderate or heavy rain can cause considerable damage even to cargo usually considered quite unperishable. The top layers of the goods may be completely soaked. Cardboard packaging may be weakened to the point of being worthless. Labels, for example, on canned goods may be loosened and even soaked off. A substantial downpour can easily cause $500 worth of damage to a truckload of unprotected merchandise; for ten trucks carrying similar merchandise, a loss of $5,000.

If a significant amount of rain is expected, the dispatcher can direct his crews to "tarp" the trucks after loading them. Tarping a large flat rack is a time-consuming operation, and the cost of doing the job, including

---

[1] The best treatment probably is J. C. Thompson and G. L. Brier, "The Economic Utility of Weather Forecasts," Monthly Weather Review, Nov. 1946.

labor time plus wear and tear on the tarping on an average trip, may amount to $20 a truck -- $200 for the fleet of trucks. To permit an early morning departure, tarping must be done the prior evening.

For the purposes of this study we consider any amount of rain in excess of .15 inches as "heavy," and assume that if rain is less than this amount no damage will result. Each evening the dispatcher must decide whether or not to tarp the trucks. Clearly it would be good policy to tarp the trucks only on evenings preceding bad days and at no other times, but in the absence of perfect forecasts this is impossible. In the absence of perfect information, the dispatcher bases his decision on the weather forecast. If the forecast is for rain, the trucks are tarped, but otherwise not. The consulting meteorologist, knowing that an unforecasted rain is extremely expensive for the trucking company, forecasts rain whenever he feels that the probability of rain exceeds a certain very low figure,[1] and thus the company very seldom is caught with unprotected trucks and a rainy day. Naturally, the forecaster often predicts rain that does not occur, and so the trucks are often protected unnecessarily. The company is happy with this arrangement and feels that the forecaster is saving them a lot of money. Let us try to calculate how much.

To calculate how much certain information is worth to a decision maker we must make some assumptions about what he would do in the absence of information. We assume throughout this paper that the decision maker always has access to records which give the historical frequency of various types of weather, or that his feel for climatological probabilities is

---

[1]This implies that the meteorologist rather than the dispatcher is really making the protect-don't protect decision. This may be undesirable in many activities, and its cause and cure will be discussed in Section IV.

reasonably accurate. If weather states tend to persist, the decision maker may be able to predict better than climatology just by looking out the window. This problem is discussed in detail in Appendix 2 of this section. But for the trucking problem, assume that, in the absence of forecasts, decisions have to be made on the basis of climatology. This means that the same decision, whatever it is, should be made every evening, for the climatological probability of rain during the day is substantially the same for a number of successive days. (Clearly this probability differs from season to season, but in Los Angeles the climatological probability of rain in excess of .15 inches is .09 during the winter rainy season, the period of our example.)

It is easy to show that if the dispatcher has to make his decision on climatology alone, he should order the trucks tarped every evening, thus incurring a daily cost of $200. The alternative, never tarping the trucks, would work out well on 91 days out of 100; but on nine days out of 100 the company would incur a loss of $5,000, an average daily loss of $450.[1]

It is useful to set down the relevant calculations in the following systematic way. Let us define the following terms:

[1] Note that we are assuming that the decision maker wishes to minimize expected costs. For the problems discussed in this report, this seems a reasonable assumption to make. In technical language, we are assuming that Von Neumann utility is proportional to money over the relevant range.

C_p = the cost of taking protective action

L = the property loss that will result if it rains and no protective action has been taken

a_1 = the decision to take protective action

a_2 = the decision <u>not</u> to take protective action

w_1 = the occurrence of rain

w_2 = the occurrence of no rain

The decision problem can be conveniently expressed in terms of the following cost matrix:

TABLE I

Weather

Decision		w_1	w_2
	a_1	C_p	C_p
	a_2	L	0

The elements of the matrix indicate what the cost to the decision maker will be if his decision is a given row, and the state of the weather turns out to be a given column. The action taken by the decision maker and the state of the weather together determine what the cost to the decision maker will be. The information given by the matrix can be written in another way:

$$c = c(a, w) , \tag{1}$$

where c or cost can be 0, C_p, or L; a can be a_1 or a_2; and w can be w_1 or w_2. In our trucking problem, $C_p = 200$ and $L = 5{,}000$.

To analyze what action the decision maker should take, we must know, in addition to his cost matrix, what information he has about the weather. Let us call climatological information, the historical frequency of the different weather states, I_o. If P_1 is the relative frequency of rainy days and $P_2 = 1 - P_1$ is the historical frequency of non-rainy days, then climatological information, I_o, is defined by the vector (P_1, P_2), where $P_1 = .09$ and $P_2 = .91$.

We have argued that if decisions must be made on the basis of climatological information alone, the decision maker should take the same action every evening. If he decides to take no protective action, decision a_2, then his average cost will be:

$$C_2 [I_o] = L P_1 + 0 P_2 = L P_1 \tag{2a}$$

$$C_2 [I_o] = 5000 \times .09 = \$450 \tag{2b}$$

The subscript "2" under C indicates action a_2. If, on the other hand, the decision maker takes protective action (decision a_1), his cost will be C_p if it rains and C_p if it doesn't rain, or his average cost will be:

$$C_1 [I_o] = C_p P_1 + C_p P_2 = C_p \tag{3a}$$

$$C_1 [I_o] = 200 . \tag{3b}$$

If, as we shall assume, the decision maker wishes to minimize expected cost, he will choose action a_1 or a_2, depending on whether C_1 or C_2 is smaller. His expected cost if he makes this "best" decision will be:

$$C^{*}[I_o] = \text{Min}\ (C_1,\ C_2) \tag{4a}$$

$$C^{*}[I_o] = \text{Min}\ (C_p,\ L\ P_1) \tag{4b}$$

$$C^{*}[I_o] = \text{Min}\ (200,\ 450) = 200\ . \tag{4c}$$

The superscript "*" over "C" indicates that the "better" action of a_1 and a_2 is taken. The "Min" of equation (4) means that C^* is the "minimum" or the smaller of the numbers C_1 and C_2.

In this example it turns out that the decision maker should always take protective action if his only information regarding the weather comes from climatology. For other problems of similar structure the best decision may be not to take protection action, but to take the occasional loss rather than the steady cost of protection. From equation 4b we can see that in the simple protect-don't protect problem the critical factor is the ratio C_p/L as compared with P_1. If C_p/L, what Thompson and Brier call the cost-loss ratio, is less than P_1, the probability of bad weather, then it pays to take protective action. But if C_p/L exceeds P_1, then the decision maker should not take protective action. Thus if the trucking company is located in an area where rain occurs only one day in fifty, then a_2 rather than a_1 would be the best choice. The loss L would then be incurred so infrequently that the costs of protection would exceed the average gains. Similarly, if the cost of protection is \$600 instead of \$200, or if the loss is \$1,000 instead of \$5,000, it would not pay to take protective action.

In our calculations of the value of weather forecasts we shall take

climatological information as a base or zero point. Our choice of this base by no means implies that climatological information, in itself, has no value. In most cases we would expect the decision maker to do worse if he did not have climatological information, but it is extremely difficult to say how much worse, and hence difficult to assess the value of climatology. The reader should be careful to interpret our value for forecast figures as value in excess of climatological information, and to keep in mind that much of the value of meteorological information may lie in the basic climatological data. We shall discuss this problem in more detail in Appendix 2 of this section.

The analysis thus far has defined the decision problem and provided a best decision and a minimum cost for the decision maker if he must make his choice among actions solely on the basis of climatology. We are now in a position to examine the question, How much are weather forecasts worth?

In the trucking example the forecast received by the decision maker is either a prediction that it will rain -- call this forecast f_1 -- or a forecast that it will not rain -- call this forecast f_2. Records show that out of 100 forecasts made during a four-month period, 18 were predictions of rain, and 82 were predictions that it would not rain. The no-rain forecasts were very reliable; on only two days did rain occur when none was forecast. This high accuracy of the f_2 forecasts was achieved because the forecaster always predicted rain whenever there was any doubt in his mind. His f_1 forecasts, therefore, were not very reliable. Indeed, on 11 of the 18 days that he forecast rain, it did not rain. It will be demonstrated later that this low skill score on the f_1 forecasts, rather than reflecting adversely on the forecaster, is an indication that he is

sensibly tailoring his forecasts to his customers' needs.

The forecasts thus define the following two-by-two contingency table:

TABLE II

Forecast

		f_1	f_2	Total
Observed	w_1	7	2	9
	w_2	11	80	91
Total		18	82	100

The entry in any column and row gives the number of days out of 100 in which a specific forecast was made, and a specific weather state occurred.

Instead of proceeding directly to find the value to the trucking company of the forecasts described in Table II, let us set up the problem more generally, find a general solution, and then substitute in the specific numbers that relate to our example.

Let:

π_1 = the relative frequency of forecast f_1 (18/100 in Table II)

$\pi_2 = 1 - \pi_1$ = the relative frequency of forecast f_2 (82/100 in Table II)

π_{11} = the conditional probability that it will rain, given that rain is forecast (7/18 in Table II)

$\pi_{21} = 1 - \pi_{11}$ = the conditional probability that it will not rain, given that rain is forecast (11/18 in Table II)

π_{12} = the conditional probability that it will rain, given that no rain is forecast (2/82 in Table II)

$\pi_{22} = 1 - \pi_{12}$ = the conditional probability that it will not rain, given that no rain is forecast[1]

[1] For the generalized contingency table:

Thus in the π_{ij}'s above, the first subscript, i, refers to the state of the weather, and the second subscript, j, refers to the forecast. Note that the π_{ij}'s are CONDITIONAL probabilities, not joint probabilities.

For our purposes, all of the information we need from Table II is contained in the vector I_f:[1]

$$I_f = (\pi_1, \pi_2; \pi_{11}, \pi_{21}; \pi_{12}, \pi_{22}) \tag{5a}$$

$$I_f = \frac{18}{100}\ \frac{82}{100}\ \frac{7}{18}\ \frac{11}{18}\ \frac{2}{82}\ \frac{80}{82}) \ . \tag{5b}$$

		Forecast f_1	f_2	Total
Observed	w_1	a	b	a+b
	w_2	c	d	c+d
Total		a+c	b+d	a+b+c+d = N

$a+c/N = \pi_1$, $b+d/N = \pi_2$
$a/a+c = \pi_{11}$, $c/a+c = \pi_{21}$, $b/b+d = \pi_{12}$, $d/b+d = \pi_{22}$
And note that if N is very large $a+b/N = P_1$, $c+d/N = P_2$.

[1] A very important relationship between I_f and I_o should be noted here. If the relative frequencies of good and bad weather over the period for which forecast data are available equal the climatological relative frequencies, then:

(a) $\pi_{11}\,\pi_1 + \pi_{12}\,\pi_2 \equiv P_1$

$$\frac{7}{18}\cdot\frac{18}{100} + \frac{2}{82}\,\frac{82}{100} \equiv \frac{9}{100}$$

(b) $\pi_{21}\,\pi_1 + \pi_{22}\,\pi_2 \equiv P_2$

$$\frac{11}{18}\cdot\frac{18}{100} + \frac{80}{82}\cdot\frac{82}{100} \equiv \frac{91}{100}$$

The first of the identities above states that the probability that it will rain and that rain is forecast plus the probability that it will rain and that rain is not forecast equal the probability of rain. The second identity states the same relationship for no rain. Thus the vector I_f is constrained by the vector I_o . And given any vector I_f , it is possible to deduce I_o .

Note that I_f is simply a shorthand notation for the bivariate distribution of forecasts and weather states.

We can now calculate how well the decision maker will do if he bases his action on information I_f . If he receives forecast f_1 , the conditional probability of rain is π_{11} . If he chooses action a_2 , his expected cost will be $\pi_{11}L$. If he chooses action a_1 , his cost will be C_p with certainty. He should choose a_1 or a_2 , depending on whether C_p or $\pi_{11}L$ is smaller:

$$C^*(f_1) = \text{Min } (C_p, \pi_{11}L) . \tag{6}$$

Similarly if he receives forecast f_2 , he should choose a_1 or a_2 , depending on whether C_p or $\pi_{12}L$ is smaller:

$$C^*(f_2) = \text{Min } (C_p, \pi_{12}L) . \tag{7}$$

Equations (6) and (7) define the decision maker's best action function or decision rule, a function that tells him the action he should take for each signal. We can write this function as:

$$a = \hat{\alpha}(f) . \tag{8}$$

If we substitute into equations (7) and (8) the appropriate numbers it is easy to see that this best decision rule, α , is: if the forecast is f_1 , take action a_1 ; and if the forecast is f_2 , take action a_2 . In Table III we have listed all possible decision rules. Our $\hat{\alpha}$ is α_2 .

TABLE III

		Forecast	
		f_1	f_2
Decision Rule	α_1	a_1	a_1
	α_2	a_1	a_2
	α_3	a_2	a_1
	α_4	a_2	a_2

We found that $\hat{\alpha} = \alpha_2$ by solving equations (6) and (7). Had we wanted, we could have started with Table III and, using the cost matrix and the information vector, could have calculated expected cost for α_1, α_2, α_3, and α_4. We would have found that using α_2 yielded the smallest expected cost. However, our indirect method of finding the best decision rule clearly saves computation.

Using equations (7) and (8), we easily can calculate average cost for the best decision rule. Since f_1 occurs with relative frequency π_1 and f_2 occurs with relative frequency $\pi_2 = 1 - \pi_1$, the decision maker's average cost will be:

$$C^*(I_f) = \pi_1 C^*(f_1) + \pi_2 C^*(f_2) \tag{9a}$$

$$C^*(I_f) = \pi_1 \text{ Min } (C_p, \pi_{11}L) + \pi_2 \text{ Min } (C_p, \pi_{12}L) . \tag{9b}$$

Notice that if we solve the equations above for the decision maker's average cost, using the best decision rule, we also solve for the best decision rule at the same time. We do not have to solve for the best

decision rule first.

Let us perform the above calculations:

$$C^*(I_f) = \frac{18}{100}\,\mathrm{Min}(200,\ \frac{7}{18} \times 5000) + \frac{82}{100}\,\mathrm{Min}(200,\ \frac{2}{82} \times 5000) \qquad (10a)$$

$$C^*(I_f) = \frac{18}{100} \times 200 + \frac{82}{100} \times \frac{2}{82} \times 5000 = 136\ . \qquad (10b)$$

Thus our decision maker, utilizing forecasts I_f , incurs an average daily cost of \$136. We may define the value of information I_f as:

$$V(I_f) = C^*(I_o) - C^*(I_f) \qquad (11a)$$

$$V(I_f) = 200 - 136 = 64\ . \qquad (11b)$$

Thus $V(I_f)$ measures how much better the decision maker will do if he bases his decision on I_f instead of I_o . Since his average daily cost will be \$200 if decisions are made on the basis of climatology, the value, over climatology, of I_f is \$64 a day.

The trucking company benefits from the forecasts it receives because the forecasts for good weather are so reliable that it can afford <u>not</u> to take protective action with no rain is forecast. This is an interesting twist to the situation many people seem to regard as the normal one, the situation in which in the absence of weather forecasts the decision maker would never take protection action. Our example shows that it is very misleading to say that forecasts have value because they permit the decision maker to take protective action. In our example forecasts have value because they permit the decision maker to take protection action less often than he would have to were there no forecasts. Or, in general, forecasts have value because they improve the decision maker's chances of doing the right thing.

It also is worth noting here that the trucking company can benefit from weather forecasts only if the accuracy of the f_2 forecasts is very high. Assume that the weather forecaster is less willing to predict rain on days when he has only the slightest fears. By predicting fair weather instead of rain on these days, he no doubt could improve his low score on f_1 or rain predictions. But he would pay a price in the accuracy of his f_2 forecasts. Assume that the forecaster could raise his score on rain predictions from 7/18 to 5/8 by predicting dry weather rather than rain on ten of the doubtful days. His average on his no-rain forecasts would then be 88/92 rather than 80/82.[1] But this decline in accuracy of the f_2 forecasts would reduce the value of the forecasting service to zero.

For if the probability of rain occurring when no rain is forecast is 4/92, then the decision maker should take protective action even when good weather is forecast. If no protective action is taken on those days, average losses would be 4/92(5000), or about $220 a day, that is, more than the cost of taking protective action. Thus if the accuracy of the f_2 forecasts is lower (the reader can easily see that the critical cutoff point is a probability of being wrong less than C_p/L , or 200/5000), the decision maker should ignore the forecasts and take protective action even on days forecast as good. His average daily cost would, of course, be $200, and the value of the forecasts, over climatology, would be zero.

Our theoretical framework can provide a clue to the value of improvements in weather forecasts. In particular, we can easily set an upper bound on the value of improved meteorological information to the decision

---

[1] There would be ten more days when good weather is predicted, and of these, eight would be dry and two rainy.

maker by asking the question, How much better would the decision maker do if he had perfect forecasts?

If forecasts were perfect, rain would be predicted on every rainy day and dry weather predicted on every dry day. Thus π_{11} , the probability of rain occurring on days forecasted as rainy, and π_{22} , the probability of good weather occurring on days forecasted as good, would both be one, and π_{12} and π_{21} would both be zero. Since forecasts clearly do not affect the relative frequencies of good and bad days, π_1 would equal P_1 and π_2 would equal P_2 ; that is, the relative frequencies of the forecasts would equal climatological probabilities. Let us call this perfect weather information I_{oo} .

$$I_{oo} = (P_1, P_2; 1, 0; 0, 1) . \tag{12}$$

If weather forecasts were perfect, then the decision maker would take protective action only on days forecast as rainy (unless $C_p > L$), and he would never incur a loss. It is obvious that since π_1 equals P_1 , $C^*(I_{oo}) = P_1 C_p$. In the trucking example, protective action would be taken on nine days out of 100, and average daily cost would thus be $.09 \times \$200 = \$18$. But let us derive $C^*(I_{oo})$ from our theoretical framework. Using equations (6), (7), and (9),

$$C^*(f_1) = \text{Min}\ (C_p,\ \pi_{11}L) = C_p \tag{13}$$

$$C^*(f_2) = \text{Min}\ (C_p,\ \pi_{12}L) = 0$$

$$C^*(I_{oo}) = \pi_1\ C^*(f_1) + \pi_2\ C^*(f_2) = P_1\ C_p . \tag{14}$$

Since average daily cost, given present meteorological information, is \$136 a day, an upper bound of \$118 a day can be placed on the value of improvements in meteorological information.

$$\text{Max } \Delta V = C^*(I_f) - C^*(I_{oo}) = 136 - 18 = 118 \; . \tag{15}$$

Any actual improvements in weather information must be worth less than this amount.

In the preceding example there were two different actions, a_1 and a_2 , two different relevant weather states, w_1 and w_2 , and two different possible forecasts, f_1 and f_2 . But the number "two" has no talismanic value. Nor is there any reason why the number of possible actions should equal the number of relevant weather states. Equally true, but less obvious, there is absolutely no reason why the number of possible forecasts need equal the number of relevant weather states.

In our preceding example, I_f was a very special type of forecast. Even though the weather states relevant from the point of view of the decision maker are rain and no rain, there is no reason why the only possible forecasts need be so limited. For example, there might be three forecasts: f_1 , almost certainly rain; f_2 , chance of rain; f_3 , almost certainly no rain. The information vector, I_f , would be

$$I_f = (\pi_1, \pi_2, \pi_3; \pi_{11}, \pi_{21}; \pi_{12}, \pi_{22}; \pi_{13}, \pi_{23}) \; , \tag{16}$$

with constraints

$$\begin{aligned} \pi_{11}\, \pi_1 + \pi_{12}\, \pi_2 + \pi_{13}\, \pi_3 &= P_1 \\ \pi_{21}\, \pi_1 + \pi_{22}\, \pi_2 + \pi_{23}\, \pi_2 &= P_2 \end{aligned} \tag{17}$$ [1]

---

[1] Of course the following are constraints:

$$\pi_1 + \pi_2 + \pi_3 = 1$$

$$\pi_{11} + \pi_{21} = 1$$

$$\pi_{12} + \pi_{22} = 1$$

$$\pi_{13} + \pi_{23} = 1$$

Or, instead of only three possible forecasts, there might be many, as many forecasts as the weather forecaster thinks he can make. However, it will be shown in Chapter IV that from the point of view of the decision maker, the number of distinctly different forecasts can be limited to the number of actions among which he must choose.

In the preceding example the alternative actions open to the decision maker were conveniently described as "protection," or "no protection." But for many problems even the notion of "protection" is misleading. For example, a trucking company can choose among alternative routes to a destination, and the choice of one route or another is not quite the same thing as taking or not taking protective action. Or consider the problem of a motion picture company with several scenes to shoot. Each scene requires a slightly different cast, and some scenes require sunny weather while others do not. The company needs weather forecasts for deciding which scene to shoot. Not only is the concept "protective action" an inept one for describing the choice, but the concept "adverse weather" has no meaning in this context, nor can the entries in the cost matrix be interpreted conveniently as a loss, or the cost of protection. The example of truck protection, though very convenient for the purposes of computation, can be a misleading one for generalization.

But a straightforward extension of the analysis is capable of handling most decision problems. The theoretical framework is that of statistical decision theory -- a hybrid from crossing statistics, psychology, and economics.[1]

---

[1] For a good survey of the field, see H. Chernoff and L. Moses, Elementary Decision Theory, Wiley, New York, 1959.

TOWARD A MORE GENERAL ANALYSIS

Assume that the decision maker must choose one particular action from some set of possible actions. Given the particular action he chooses, a_i , and the actual state of nature (for our purposes the weather), w_j , he receives a certain net revenue, $r(a_i, w_j) = r_{ij}$, or incurs a certain cost $c(a_i, w_j) = c_{ij}$. If the decision maker knew what the weather was going to be, he would be able to pick his best action. But the decision maker does not generally know this and so must make his decision on the basis of incomplete information, say, a particular forecast, f_k . It is convenient to assume that the decision maker knows the conditional probabilities of the different weather states (the π_{jk}'s).[1] Thus given any particular f_k , the decision maker can choose the action which minimizes expected cost.

The joint distribution of forecasts and actual states of the weather, $P(w, f)$, defines the decision maker's information vector, I_f . Given the cost matrix $c(a, w)$ and a particular I_f , we can calculate for each forecast the action the decision maker should take if he wishes to minimize expected cost. Thus we can find the $\hat{\alpha}$ which, for a given I_f , minimizes $E(c) = g(\alpha, I_f)$, where $\hat{\alpha}$ is the optimum "action function," "strategy," or "decision rule," $a = \hat{\alpha}(f)$. Given a particular I_f , we can also calculate (although the decision maker need not) the relative frequencies of the different forecasts. Then we are able, in a straightforward way, to find out how well the decision maker will do on the average with different

---

[1] Thus our analysis can be strictly Bayesian. Instead of dealing with the probability distribution of signals (forecasts) for each actual state of the weather, and with the a priori distribution of states of the weather, we find it more convenient to deal with the probability distribution of different weather states for each signal (forecast), that is, deal directly with the posteriori probabilities, and with the probability distribution of forecasts.

information vectors, I_f . Thus we are able to value different information vectors, including the existing one, against each other and against climatology. The general computational framework is presented in Appendix 1 to this section. In Appendix 2 we briefly examine some issues relating to "climatology" and "persistence."

Little of the work in applied meteorology has been cast within the framework outlined above, but there are a few theoretical studies which have set a precedent for this type of research. J. C. Thompson has written several papers in which he presents a cost matrix and attempts to value weather forecasts.[1] Irving Gringorten and T. A. Gleeson have written papers presenting a similar analysis.[2] Empirical research along the above lines has been very meager. Although there have been several excellent studies undertaken of the effect of weather on various sorts of human activities, only Thompson has attempted to value forecasts within the theoretical framework outlined above. It is hoped that our paper will stimulate further research in this direction.

In Section II we present several case studies cast within this framework. The reader who is not interested in applications of the theory may go on directly to Section III.

The simple theoretical framework we have given above does not fit all

---

[1] "The Economic Utility of Weather Forecasts," Monthly Weather Review, Nov. 1946. "On the Operational Deficiencies of Weather Forecasts," Bulletin of The American Meteorological Society, June 1952, pp. 1-9.

[2] I. Gringorten, "Forecasting by Statistical Inferences," Journal of Meteorology, Vol. 7, No. 6, and Probability Estimates of the Weather in Relation to Operational Decisions, Geophysics Research Directorate, Air Force Cambridge Research Center. T. A. Gleeson, "A Production and Decision Method for Applied Meteorology and Climatology Based Partly on the Theory of Games," Journal of Meteorology, Vol. 17, April 1960.

decision problems equally well. Certain problems, particularly problems in which decisions are made sequentially, are more usefully treated within a slightly different framework. In section III we study a special class of problems, which we call "delay" problems, and develop an appropriate special-purpose framework.

The discussion thus far has been from the point of view of a decision maker with a given cost matrix whose information about the weather is summarized by a given information vector. In Section IV we turn our attention to the forecaster. The weather forecaster is responsible for translating observations into forecasts, a process that generates the decision maker's I_f. While the decision maker is interested only in those aspects of the weather which enter his cost matrix, the forecaster is interested in meteorological information which permits him to make good forecasts. We shall assume that the forecaster may observe the w's relevant to the decision maker's costs (although he observes what the variables <u>are</u> and the decision maker is interested in what they <u>will be</u>), other meteorological variables z, and of course knows the relevant date and location, t and L. We shall model the forecaster's problem as that of choosing an optimum η for the equation $f = \eta(w, z, t, L)$, where f is the forecast made, and w, z, t, and L are the data the forecaster works with. It should be noted that a given η generates a given I_f.

Even less work has been done in applied meteorology on the forecaster's problem, though Thompson has studied the problem in some detail,[1] and several of the references cited earlier do deal with certain aspects of it.

---

[1] J. C. Thompson, "A Numerical Method for Forecasting Rainfall in the Los Angeles Area," <u>Monthly Weather Review</u>, July 1950.

II. SELECTED CASE STUDIES

INTRODUCTION

In this Section we examine several empirical decision problems to which the theoretical framework of Section I can be applied in a straightforward way. No new theory will be presented. Despite some simplifying assumptions, we believe our models of the problems are reasonably realistic, and our numbers reasonably close to the correct figures. The examples were selected because the data were readily available, not for their social significance. They are intended to demonstrate the use of the theoretical framework, and as garden-variety examples of the multitude of human activities in which weather information has value.

Before we examine the case studies, it seems important to spell out a few of their limitations. The empirical data required by the theory are the cost matrix, $c = c(a, w)$, and the information vector, I_f. If there are n possible actions and m relevant states of the weather, the cost matrix will be n by m. The information vector is equivalent to a contingency table, of course; if there are q possible forecasts the contingency table will be q by m. $V(I_f)$ can be computed directly from $c(a, w)$ and I_f. Thus the theory is straightforward; obtaining the data to fit the theory, or even adjusting the theory to fit a particular case, is a much tougher nut to crack.

First, consider the problem of defining the possible alternative actions in any decision problem. In the example of Section I, a unique protective action was assumed; but even in very simple problems, there are likely to be several alternatives and certainly several different intensities

of protective action. If the decision maker is the manager of a municipal snow-clearance agency, he can alert trucks and men and disperse them to appropriate locations, check sand reserves, alert a standby force, and so on. He need not take all these actions at the same time, of course, and he can vary their completeness. If there is a low probability of snow, for example, he may decide merely to check sand supplies and send trucks and men to cover one or two major streets. In virtually every case, then, the analyst faces a difficult problem in exhaustively listing a finite and reasonably small set of mutually exclusive actions the decision maker might take.

It is likely to be even more difficult to list all states of the weather that are relevant to the decision problem. In the example of Section I, the decision maker's only concern was whether more than .15 inches of rain would fall. Otherwise, the specific amount of rain, the temperature, visibility, etc., were irrelevant. But in the snow-alert problem above, the volume of snow is decidedly important, and visibility and temperature are also likely to be, since even a little snow can do a great deal of mischief if it creates icy surfaces and makes visibility poor. Thus, the snow-clearance manager's decision to take a state of alert should probably depend on temperature and visibility forecasts as well as on forecasted amounts of snowfall. For a farmer deciding what crop to plant and when, rainfall and its time distribution, temperature and its time distribution, and many other weather factors are likely to be important. To analyze the problem properly, all relevant weather conditions and dimensions must somehow be grouped into a finite set of relevant weather states.

Of course the problems of listing a_i and w_j are inseparably related

to the problem of calculating c_{ij} , the cost if action i is taken in weather state j . The longer the list of a_i and w_j , the more accurately can c_{ij} be calculated, but the longer and more complex will be the study. If on the other hand the listings of a_i and w_j are kept small, each c_{ij} must be considered as the mean of a frequency distribution which may have a large variance.

Frequently, it is not easy to quantify all the c_{ij}'s . For example, weather information is clearly valuable to newspaper distributors, but it is very difficult to calculate how valuable. The problem is not the fact that there are many different possible actions and relevant weather states; it is mainly one of rain or no rain. Almost any amount of rain will dampen a folded newspaper enough to cause customer dissatisfaction. If any rain is expected, the people responsible for proper delivery generally require that the papers be "waxed" (protected with a cover of waxpaper). This protective action involves a slight added cost for waxpaper and increased folding time. This cost is known fairly accurately: about $1.50 for 250 papers. If the papers are not waxed and they get wet, loss is sustained but it is difficult to calculate it. Probably the major loss is social: the decreased utility of the newspaper to its readers. Some of this loss may be passed back to the distributor, but policies of the independent dealers vary somewhat as to extra deliveries, etc. Usually, extra deliveries are not made to replace wet papers with dry ones, and monthly subscription rates are not adjusted for ruined papers. Therefore, neither the distributor nor the paper carrier suffers a direct financial loss, but there is certainly a loss of goodwill or customer satisfaction, which may result in cancelled subscriptions.

The problem of obtaining the I_f's is also sometimes troublesome. Since the format and vocabulary of weather forecasts may differ from locality to locality and even from person to person, it is often difficult to categorize them. Furthermore, if the forecasts are for specially selected localities, it may be difficult to develop an accurate contingency table, that is, to calculate the relative frequencies of the different weather states for each type of forecast; and often, even the relevant climatological data are difficult to obtain.

The difficulties mentioned above do not imply that empirical research within the framework of the theory is impossible; the case studies presented in this paper are evidence to the contrary. The implication is rather that such empirical studies are difficult, and that specific numbers should be taken with a grain of salt.

THE VALUE OF METEOROLOGICAL INFORMATION TO THE MOTION PICTURE INDUSTRY

The filming of motion pictures is extremely sensitive to the weather. This is the primary reason the industry originally concentrated in Southern California; but even in this generally favorable climate, producers quickly saw that weather information could be extremely helpful in keeping costs down.

But before turning to the use of shorter-range forecasts, we should note that climatological information plays a valuable role in many of the decisions that must be made. Although most pictures are filmed locally in Southern California, either at the studio or at nearby locations, many others are filmed at distant locations throughout the country and the world. Climatological data provide producers with information on the most favorable time of the year to shoot in different locations, the expected number of

days which will be lost because of bad weather, and other relevant matters. Even for pictures filmed locally, climatological data enter the decision as to the best time of the year to shoot.

Once a schedule has been tentatively set and actual operations have begun, the production manager is faced with day-to-day scheduling decisions. Should he plan to shoot an indoor scene or an outdoor scene? If outdoor, should he bring special lighting equipment for protection in case the day turns out to be dark? For such decisions, good weather information often has great value.

A Protection Problem

For example, consider the problem of a motion-picture producer filming an outdoor TV series in color. It is early summer, when there is almost no chance of rain in the coastal Southwest area where the pictures are being taken, but the weather must also be sunny if no special lighting equipment is at hand. The cost of hiring special lighting crews and equipment may amount to approximately \$1,000 a day. They will not be used if the weather is sunny; but if it is cloudy and they are not available, about \$10,000 in salaries, equipment rentals, etc. will be lost. On the average it is cloudy enough to hamper operations about two days a month. Thus, $I_o = (2/30, 28/30)$, and the cost matrix can be written:

	w_1	w_2
a_1	1000	1000
a_2	10,000	0

Thus:

$$C^*(I_o) = \text{Min } (1000,\ 2/30 \times 10{,}000) = 667 \ .$$

In the absence of weather forecasts, the decision maker's best action is to dispense with special lighting crews and take his losses when they occur.

The producer, realizing the value of good information, receives daily forecasts from a consulting meteorologist. The forecasts are either for cloudy or clear weather:

$$I_f = (\pi_1, \pi_2; \pi_{11}, \pi_{21}, \pi_{12}, \pi_{22}) = (\frac{3}{30}, \frac{27}{30}, \frac{1}{3}, \frac{2}{3}, \frac{1}{27}, \frac{26}{27}) .$$

Calculating $C^*(I_f)$ and $V^*(I_f)$:

$$C^*(I_f) = \frac{3}{30} \text{Min}[1000, \frac{1}{3}(10,000)] + \frac{27}{30} \text{Min}[1000, \frac{1}{27}(10,000)] = 433 ;$$

and

$$V(I_f) = 234 .$$

Thus if the decision maker hires a special lighting crew to go out on days forecast as bad, he will make, on the average, \$234 more per day than he will if he has to base his decisions on climatology.

How much would it benefit our decision maker if he had perfect forecasts? This figure will put an upper bound on the value of improved meteorological information for this decision maker. Assume $\pi_{11} = 1$, $\pi_{22} = 1$, $\pi_1 = P_1$ and $\pi_2 = P_2$; then

$$C^*(I_{oo}) = P_1 \text{ Min } (1000, 10,000) + P_2 \text{ Min } (1000, 0) = \frac{2}{30} \times (1000) = \$66 .$$

Since with present forecasts average costs are \$433, improved forecasts may be worth \$433 - \$66 = \$367 a day, as an upper bound.

The Two-Way Call

In the case we have just examined the threat of rain was so small it could be ignored. But at different times of the year and at different locations the threat of rain may be of major importance to the decision maker, particularly if the film is to have both indoor and outdoor scenes. Preparations for shooting a scene usually must be made at least a day in advance. The cast must be alerted, equipment and props rented, lunches and transportation arranged. If the production manager plans to shoot a scene outdoors and the weather turns rainy he may lose as much as $15,000 to $20,000. If he shoots indoors when the weather is good he loses nothing but a day of good weather. But if there is a subsequent run of bad weather and only a small number of scenes can be shot indoors, the waste of a good day may be costly because it has delayed completion of the picture.

For these reasons, production managers usually try to complete a good share of their outdoor scenes early. Although it is extremely difficult to assign a number to the value of shooting an outdoor scene rather than an indoor scene, producers behave as if this number were quite large.[1] Let us assume that at the current stage of progress of a particular picture, the production manager feels it is worth about $9,000 more to be able to complete an outdoor scene than an indoor scene. This figure is net of costs. Lighting equipment is included in the costs of the outdoor scene; the production manager has decided to bring it along for protection whenever he tries to shoot an outdoor scene. If it rains at all the outdoor scene cannot be

---

[1] In principle, it is possible to calculate this number; in the following sections we shall show several of the necessary steps. It turns out that our result, $V(I_f)$, is insensitive to changes in the $9,000 figure.

filmed. During February the climatological probability of rain is 0.17. The production manager must balance the costs of planning to shoot outdoors and getting rained on, against the benefits of completing an outdoor scene -- the costs and benefits, of course, weighted by the relevent probabilities.

On the night before shooting, the production manager can take one of three actions: he can call up and firmly commit to hire the indoor people, do the same for the outdoor people, or place what is known as a "two-way call." Before describing the two-way call, let us examine the first two possible actions.

If the production manager firmly hires the indoor people, action a_1, he will shoot indoors regardless of the weather; let us say the net value of this action is zero.[1] If the production manager firmly hires the outdoor people, action a_2, and the weather turns out to be good, he will do \$9,000 better. If it rains, however, he is in trouble; he is committed to pay full salaries to the outdoor people, but he cannot shoot outdoors. However, he can salvage something, but not very much, by using outdoor people, and people who are part of both the indoor and outdoor casts, to shoot indoor scenes which require no advance preparation. He values this possibility as \$16,000 less than a planned indoor shot -- a figure close to but not quite equal to his expenses of the day.

	w_1	w_2
a_1	0	0
a_2	-16,000	9,000

[1] Since everything is measured as differences from the value of shooting indoors, there is no loss of generality and some gain in computational convenience by setting this value equal to zero.

The third alternative, a two-way call, is a hedging operation involving minimal preparation for two work programs the following day -- one for indoors and one for outdoors. Then, on the basis of a very reliable[1] short-range forecast made the next morning, the production manager may decide which scene to shoot. Extra costs are involved in making a two-way call, however. If the company goes outdoors, some of the cast, electricians, etc., who were hired for indoor work must be cancelled at a cost of about \$4,000. If the company goes indoors, some of the outdoor cast, lunches, transportation, etc. must be cancelled at a cost of about \$3,000.

Considering all three possible actions, we can write the net revenue matrix:

	w_1	w_2
a_1	0	0
a_2	-16,000	9,000
a_3	- 3,000	5,000

Since the decision maker wishes to choose the action which will maximize his expected net revenue, this is a problem of maximization rather than minimization. It has already been stated that climatological probabilities are $I_o = (.17, .83)$. Calculating expected net revenue for each action:

$$R_1(I_o) = 0 ,$$

$$R_2(I_o) = .17(-16,000) + .83(9000) = -2720 + 7470 = 4750 ,$$

$$R_3(I_o) = .17(-3000) + .83(5000) = -510 + 4150 = 3640 , \text{ and}$$

$$R^*(I_o) = 4750 .$$

---

[1] For our computations we will assume the forecast is perfect.

Note that R^* is the largest of R_1, R_2, and R_3. In the absence of any weather forecasts, the production manager should commit himself tonight to hiring the outdoor cast for tomorrow, with an expected revenue of \$4,750.

The forecasts the motion-picture company receives are in three possible categories. Forecast f_1 is a high-confidence forecast of rain made a day in advance. It is received about one day in ten, and is accurate .67 per cent of the time. Thus $\pi_1 = .1$, $\pi_{11} = .67$, and $\pi_{21} = .33$. Forecast f_3 is a high-confidence good weather forecast in which $\pi_3 = .70$, $\pi_{13} = .04$, and $\pi_{23} = .96$. Forecast f_2 is in between -- a forecast of uncertain weather in which $\pi_2 = .2$, $\pi_{12} = .4$, and $\pi_{22} = .6$.[1] Thus

$$I_f = (.1,\ .2,\ .7;\ .67,\ .33;\ .4,\ .6;\ .04,\ .96)\ .$$

If the forecast is f_1, the production direction should plan to shoot indoors, as shown below:

$$R^*(f_1) = \mathrm{Max}[0,\ .67(-16{,}000) + .33(9000),\ .67(-3000) + .33(5000)] = 0\ .$$

Similarly, it can be shown that when the forecast is f_2, the production manager should put in a two-way call, and when it is f_3, he should plan to shoot outdoors:

$$R^*(f_2) = \mathrm{Max}[0,\ .4(-16{,}000) + .6(9000),\ .4(-3000) + .6(5000)] = 1800\ ;$$

$$R^*(f_2) = \mathrm{Max}[0,\ .04(-16{,}000) + .96(9000),\ .04(-3000) + .96(5000)] = 8000\ .$$

If he makes the correct decisions, the production manager's expected revenue is:

---

[1] Note that

$$\pi_{11}\pi_1 + \pi_{12}\pi_2 + \pi_{13}\pi_3 \cong P_1$$

$$.67(.1) + .40(.20) + .04(.7) \cong .17\ .$$

The slight discrepancy is due to rounding.

$$R^*(I_f) = \pi_1 R^*(f_1) + \pi_2 R^*(f_2) + \pi_3 R^*(f_3) = 0 + .2(1800) + .7(8000) = 5960 \ ,$$

And the value of the forecast is:

$$V^*(I_f) = 5960 - 4750 = 1210 \ .$$

To put an upper bound on the possible value of improved weather forecasts, consider how well the decision maker would do if he had perfect weather forecasts, $R^*(I_{oo})$: [1]

$$R^*(I_{oo}) = .17 \ \mathrm{Max}(0, -16{,}000, -4000) + .83 \ \mathrm{Max}(0, 9000, 5000) = 7470 \ .$$

Since with present forecasts the decision maker can make an expected profit of 5,960, an upper bound on the value of any improvements in forecasts is 7,470 - 5,960 = 1,510.

A Sensitivity Analysis of the Two-Way Call

Little confidence can be reposed in the $9,000 figure we assigned to the value of shooting outdoors as opposed to indoors. It therefore seems worthwhile to replace it with X, and see how sensitive $V(I_f)$ is to differing values of X. It turns out that our choice of $X = \$9000$ is not critical. Any value of $X > \$2000$ yields similar values of $V(I_f)$.

Rewriting the cost matrix as a function of X, and rewriting the information vector for easy reference, we have:

---

[1] The perfect forecasts considered here are evening forecasts of the next day's weather. They are not to be confused with the perfect forecast we assume is made the next morning in our discussion of the two-way call.

	w_1	w_2
a_1	0	0
a_2	-16,000	X
a_3	- 3,000	X -4,000

$$I_f = (.1,\ .2,\ .7;\ .67,\ .33;\ .4,\ .6;\ .04,\ .96)\ .$$

Rewriting $R^*(I_o)$ as a function of X, we have:

$$R^*(I_o) = \text{Max}[0,\ .17(-16{,}000) + .83(X),\ .17(-3000) + .83(X - 4000)] \qquad (1)$$

$$= \text{Max}[0,\ -2720 + .83(X),\ -3830 + .83(X)]$$

$$= \begin{cases} 0;\ X < 3277 \text{ (action } a_1) \\ .83(X) - 2720;\ X \geq 3277 \text{ (action } a_2)\ . \end{cases}$$

If the decision maker has only climatological information, his best action is a_1 if $X < 3277$ and a_2 if $X \geq 3277$. He never should use a_3.

Rewriting $R^*(f_1)$, $R^*(f_2)$, and $R^*(f_3)$ as a function of X yields;

$$R^*(f_1) = \text{Max}[0,\ .67(-16{,}000) + .33(X)\ .67(-3000) + .33(X - 4000)] \qquad (2)$$

$$= \text{Max}[0,\ -10{,}720 + .33(X),\ -3330 + .33(X)]$$

$$= \begin{cases} 0;\ X < 10{,}090 \text{ (action } a_1) \\ .33(X) - 3330;\ X \geq 10{,}090 \text{ (action } a_3)\ ; \end{cases}$$

$$R^*(f_2) = \text{Max}[0,\ .4(-16{,}000) + .6(X),\ .4(-3000) + .6(X - 4000)] \qquad (3)$$

$$= \text{Max}[0,\ -6400 + .6(X),\ -3600 + .6(X)$$

$$= \begin{cases} 0;\ X < 6000 \text{ (action } a_1) \\ .6(X) - 3600;\ X \geq 6000 \text{ (action } a_3)\ ; \end{cases}$$

$$R^*(f_3) = \text{Max}[0,\ .04(-16{,}000) + .96(X),\ .04(-3000) + .96(X - 4000)] \quad (4)$$

$$= \text{Max}[0,\ -640 + .96(X),\ -3960 + .96(X)]$$

$$0;\ X < 666 \text{ (action } a_1)$$

$$.96(X) - 640;\ X \geq 666 \text{ (action } a_2)\ .$$

Let us plot $R^*(I_o)$, $R^*(I_f)$, and $V(I_f)$ as a function of X (see Figure 1). If $X < 666$, then:

(1) $R^*(I_o) = 0$

(2) $R^*(I_f) = 0$

(3) $V(I_f) = 0$.

If $666 \leq X < 3277$, then:

(1) $R^*(I_o) = 0$

(2) $R^*(I_f) = .7[.96(X) - 640] = .672(X) - 448$

(3) $V(I_f) = .672(X) - 448$.

If $3277 \leq X < 6000$:

(1) $R^*(I_o) = .83(X) - 2720$

(2) $R^*(I_f) = .672(X) - 448$

(3) $V(I_f) = 2272 - .158\,(X)$

If $6000 \leq X < 10{,}900$:

(1) $R^*(I_o) = .83(X) - 2720$

(2) $R^*(I_f) = [.672(X) - 448] + .2[.6(X) - 3600] = .792(X) - 1168$

(3) $V(I_f) = 1552 - .038(X)$.

If $10{,}900 \leq X$:

(1) $R^*(I_o) = .83(X) - 2720$

(2) $R^*(I_f) = [.792(X) - 1168] + .1[.33(X) - 3300] \approx .83(X) - 1498$ [1]

(3) $V^*(I_f) = 1222$.

---

[1] The slight discrepancy is due to a rounding error.

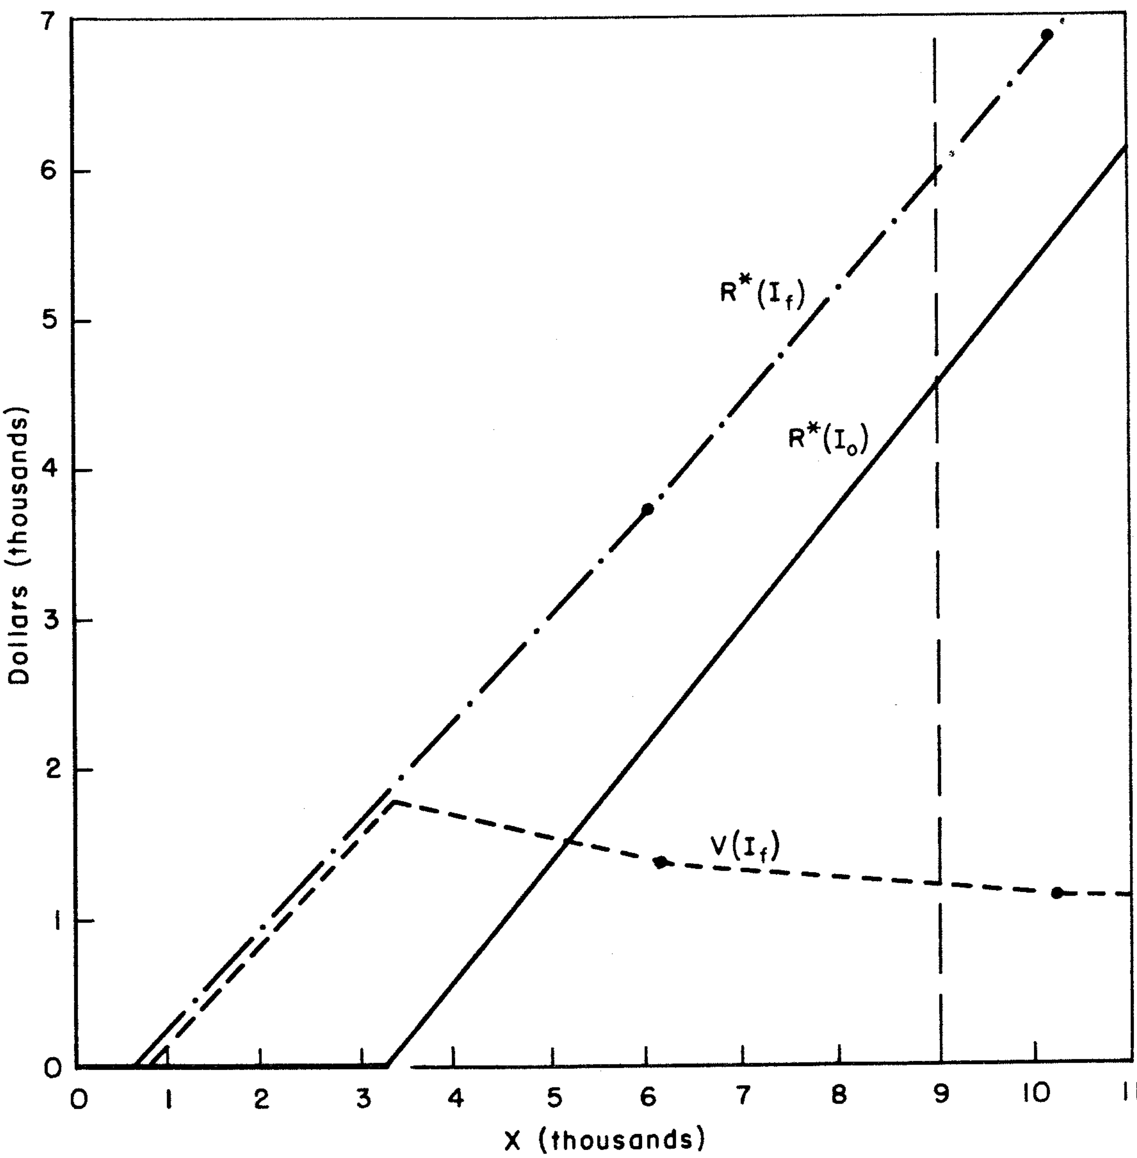
7
6
5
4
3
2
1
0
Dollars (thousands)
$R^*(I_f)$
$R^*(I_o)$
$V(I_f)$
0
1
2
3
4
5
6
7
8
9
10
11
X (thousands)

Fig. 1

Figure 1 shows that $V(I_f)$ is insensitive to different values of X.

THE VALUE OF METEOROLOGICAL INFORMATION TO THE CONCRETE LAYING BUSINESS

The cost of concrete slabs is sensitive to the state of the weather. It costs about $.25 a square foot to pour a slab when the ground is dry and the weather good, but considerably more when the ground is muddy. The equipment is more difficult to move, workmen have trouble with mud sticking to their tools, and in general the pace of work is slower. The situation is even worse if it rains while the slab is poured or before the slab is dry. A light rain -- .15 inches or less -- will usually mar only the surface, but a moderate or heavy rain can ruin the slab.

Let us focus on the rain-damage problem. If a moderate or heavy rain falls while the slab is poured, or shortly thereafter, the mixture may be diluted. If a core test taken later reveals the strength of the slab to be sub-standard, the slab will have to be ripped out and another pour made. Breaking up the concrete, loading it on trucks, and transporting it from the site is a very costly operation. If the contractor decides, while the concrete is still fresh, that his slab is likely to fail subsequent tests, he may remove it at once and repour at a later date. The removal and re-pour together cost about $.50 a square foot -- somewhat less than when the concrete has hardened.

If a light rain falls during the final trowelling process, the contractor can save the slab by applying a thin top mix the following day. Doing so produces a fairly smooth but slightly sandier surface. The additional cost for this step is about $.03 a square foot. If a light rain is falling, a wet surface can also be protected by covering it with tarps, newspaper, etc., at a slightly greater labor cost.

Let us simplify the problem somewhat. Consider a contractor who has many slabs to pour, can pour one a day, and who pays no penalty for taking a long time to complete a job. Each day he must decide whether or not to try to pour a slab. Assume that if he does not try, his cost is zero. If he completes the job he nets \$1,500, if no rain mars the surface. If he pours the slab and there is light rain he can protect and repair the surface for an added cost of about \$300. If he pours the slab and there is heavy rain, all his work is ruined -- he must rip up the slab and incur a cost of \$2,500.

Defining a_1 as the decision to pour and a_2 as the decision not to pour, w_1 as good weather, w_2 as light rain, and w_3 as heavy rain, the net revenue matrix is:

	w_1	w_2	w_3
a_1	0	0	0
a_2	1,500	1,200	-2,500

The climatological probabilities of w_1 , w_2 and w_3 are $P_1 = .84$, $P_2 = .08$, $P_3 = .08$. A quick glance at the matrix gives the same results as computation. In the absence of weather forecasts the contractor should try to pour. His average revenue would be:

$$R^*(I_o) = .84 \times 1500 + .08 \times 1200 + .08 \times -2500 = 1156 .$$

The forecasts the contractor receives are predictions either of rain or of no rain. For a one-hundred-day period the contingency table was as follows:

	f_1	f_2	Total
w_1	79	5	84
w_2	3	5	8
w_3	1	7	8
Total	83	17	100

(The numbers have been slightly adjusted to equal climatology.)

A glance at the table again gives the same results as computation; the decision maker should not try to pour on days when rain is forecast. Receiving these forecasts and deciding accordingly, the contractor will make as an average net revenue:

$$R^*(I_f) = \frac{83}{100}\left[\frac{79}{83} \times 1500 + \frac{3}{83} \times 1200 + \frac{1}{83} \times -2500\right] + \frac{17}{100} \times 0 = 1196 .$$

The value of the forecasts to him thus must be:

$$V(I_f) = R^*(I_f) - R^*(I_o) = 40 .$$

If forecasts were perfect, the contractor would never try to pour on days of heavy rain, but would pour slabs on all other days. His average net revenue would be:

$$R^*(I_{oo}) = .84 \times 1500 + .08 \times 1200 = 1356 .$$

Thus 160 per day gives an upper bound to the worth of improved weather forecasts to the contractor.

III. WEATHER FORECASTS AND ACTIONS WHICH CAN BE DELAYED

INTRODUCTION

Let us reconsider the problem facing a municipal snow-clearance manager who must base his actions on snow forecasts.[1] Snowplow equipment must be readied before it is used, and sand supplies must be stocked before they can be distributed, but the decision to have these done by no means implies that the decision maker must actually use the equipment or spread the sand. He can make that decision when he has more reliable weather information.

Many other important problems are marked by a similar sequential structure. Cranberry and citrus protection, and the hurricane-warning problem, are prime examples. To deal with them, we must develop a framework of analysis which explicitly treats decisions made over a period of time. The tools of analysis developed in this Section will be helpful in dealing with these sequential problems, and are useful in their own right.

In the problems treated here, the decision maker has two choices on any given day: to attempt or not attempt to do a given job. We shall assume only two relevant weather states: one in which the job can be done, one in which it cannot. The problem is easily treated in a straightforward manner if we assume that the same go no-go choice faces the decision maker day after day, and that undone work does not accumulate (the books are wiped clean at the end of each day). The slab-pouring problem of Section II is a case in point. The analysis can deal with daily decisions, and the calculation of average daily cost is what is relevant.

---

[1]See J. C. Thompson, "The Snowfall Probability Factor," The American City, December 1959.

But what if the concrete contractor is trying to decide whether or not to take a certain contract? He would then be concerned with the expected cost of completing the job, not merely with the average daily cost. And what if, when he commits himself to do a job, he becomes saddled with daily expenses which he cannot avoid by not working? He may have to rent special equipment for the duration of the job, for example, or pay his labor partial wages even on days when no attempt is made to pour concrete. Or, what if the contract penalizes delay in completing the job or failure to complete it by a specified date? Clearly, the expected cost of fulfilling such a contract is a function of the expected number of days that will pass before the job is done.

This kind of a problem can be handled within the general framework of Section I, but the reader may verify that the analysis can become quite complex. This Section develops some special-purpose tools and examines a few empirical cases.

A LINEAR POSTPONEMENT MODEL

Assume that the cost of completing a job can be broken down into the following three elements: the direct cost C_D of doing the work, a fixed cost or penalty charge C_F for each day that elapses before the job is completed, and an added loss C_L incurred each day the decision maker tries to do the job and the weather turns out to be unfavorable. Thus, in our slab-pouring problem, the contractor may wish to pour as soon as possible, to avoid the cost of extra days. But if he pours the concrete and it rains, he will have to rip out his work and try again. Or a film producer may wish to shoot an outdoor scene in a situation where delay costs money, but if he

attempts to do so and the day is bad, he will incur a great deal of expense and fail to do the job. The cost function is:

$$C = C_D + N_T C_F + N_L C_L \quad . \tag{1}$$

N_T is the total number of days to complete the job, and N_L is the number of days the job is attempted and there is failure because of bad weather. How does meteorological information affect the expected values of N_T and N_L?

To start, consider a special case in which the computations are reasonably straightforward. Assume that the weather is not marked by any particular persistence phenomenon, and that the climatological probabilities of bad weather and good weather are P_1 and $P_2 = (1 - P_1)$ respectively. If the decision maker is committed to do the job, then, in the absence of any weather information save climatology, his only rational strategy is to go out every day until he hits good weather and can complete the job. How many days must he expect to do this before he succeeds? Once this question is answered it is easy to calculate the expected total number of days which will elapse before the job is completed. Since the decision maker goes out on every day and fails on all but the last, $N_T = N_L + 1$.

If he goes out every day until he succeeds, the probability of success on the first day -- that $N_L = 0$ -- is P_2. The probability of success on the second day -- that $N_L = 1$ -- is $P_1 P_2$. And the probability that he will succeed on the $(K + 1)$th day, that $N_L = K$, is $P_1{}^K P_2$. Thus:

$$\bar{N}_L = 0P_2 + 1P_1P_2 + 2P_1^{\,2}\,P_2 + \ldots + KP_1^{\,K}P_2 + \ldots$$

$$= P_2\ [P_1 + 2P_1^{\,2} + 3P_1^{\,3} + \ldots + KP_1^{\,K} + \ldots\ ]$$

$$= \frac{P_2}{1-P_1}\ [P_1 + P_1^{\,2} + P_1^{\,3} + \ldots] \tag{2}$$

$$= \frac{P_2}{1-P_1}\ [\frac{P_1}{1-P_1}] = \frac{P_1}{P_2} \quad ; \qquad \text{and}$$

$$\bar{N}_T = \frac{P_1}{P_2} + 1 = \frac{1}{P_2} \quad . \tag{3}$$

Thus if the probability that the weather will be favorable is $\frac{1}{4}$, then the decision maker will, on the average, go out three days and fail before he finally succeeds. So the job will take, on the average, four days to complete.

Now, suppose the decision maker receives weather forecasts which are good enough that it pays to use them (the conditions for this assumption will be developed in the text). That is, assume he will go out only if good weather is forecast. Let π_2 be the relative frequency of good-weather forecasts and $\pi_1 = 1 - \pi_2$ be the relative frequency of bad-weather forecasts. Let π_{22} be the conditional probability of good weather, given that it is forecast, and $\pi_{12} = 1 - \pi_{22}$ be the conditional probability of bad weather, given that good weather is forecast. π_{11} and π_{21} are the similarly defined conditional probabilities for bad weather and good weather, given that bad weather is predicted.[1]

---

1 $\pi_{11}\,\pi_1 + \pi_{12}\,\pi_2 = P_1$ and $\pi_{21}\,\pi_1 + \pi_{22}\,\pi_2 = P_2$

and therefore

$$\frac{\pi_1}{\pi_2} = \frac{P_1\,\pi_{22} - P_2\,\pi_{12}}{P_2\,\pi_{11} - P_1\,\pi_{21}} \quad .$$

If the decision maker goes out only if good weather is forecast, then the probability that he will succeed the first day he goes out -- that $N_L = 0$ -- is π_{22}. The probability that he will succeed on the second day -- that $N_L = 1$ -- is $\pi_{12}\,\pi_{22}$. And the probability that he will succeed on the $(K + 1)$th day -- that $N_L = K$ -- is $\pi_{12}{}^{K}\pi_{22}$. By exactly the same steps that led to Equation (2):

$$N_L = \frac{\pi_{12}}{\pi_{22}} \quad . \tag{4}$$

If the forecasts are any good, π_{12}/π_{22}, the ratio of the probability of bad weather to that of good weather, given that good weather is forecast, should be less than P_1/P_2. Thus, the better the forecast the fewer days the decision maker will go out and experience bad weather and suffer a loss C_L. If weather forecasts are perfect, then $\pi_{12} = 0$ and $\bar{N}_L = 0$. If weather forecasts are no better than climatology (if they are worse they will not be used), $\pi_{12} = P_1$, $\pi_{22} = P_2$, and $N_L = P_1/P_2$. The decision maker will go out and suffer losses on just as many days, on the average, when he goes out when the forecast is favorable, as he would if he went out every day, regardless of the forecast.

The total number of days it takes, on the average, to complete the job equals $\bar{N}_L + 1$ (the number of days good weather is forecast) plus the number of days bad weather is forecast and the decision maker stays home. It can be shown that if there is no persistence in weather states, for

every $\bar{N}_L + 1$ days when good weather is forecast there will be, on the average, $\frac{\pi_1}{\pi_2}$ $(N_L + 1)$ days when bad weather is forecast. Since $N_L + 1 = \frac{1}{\pi_{22}}$:

$$\bar{N}_T = \bar{N}_L + 1 + \frac{\pi_1}{\pi_2} (N_L + 1) = \frac{1}{\pi_2 \pi_{22}} \quad . \tag{5}$$

If the decision maker has no forecasts to guide him, let us write the average cost as $C(I_o)$, and as $C(I_f)$ if he does have them. From (1), (2), and (3):

$$C(I_o) = C_D + C_F \frac{1}{P_2} + C_L \frac{P_1}{P_2} \quad . \tag{6}$$

And from (1), (4), and (5):

$$C(I_f) = C_D + C_F \frac{1}{\pi_2 \pi_{22}} + C_L \frac{\pi_{12}}{\pi_{22}} \quad . \tag{7}$$

Defining, as before, $V(I_f) = C(I_o) - C(I_f)$ as the value of the meteorological information, we have:

$$V(I_f) = C_F \left(\frac{1}{P_2} - \frac{1}{\pi_2 \pi_{22}} \right) + C_L \left(\frac{P_1}{P_2} - \frac{\pi_{12}}{\pi_{22}} \right) \quad . \tag{8}$$

The parameters of Equation (8) are all that are needed to calculate the value of the weather forecast I_f. If $V(I_f)$ turns out to be negative, it is better for the decision maker to ignore the forecasts and go out every day until the job is done. And $V(I_f)$, as calculated in Equation (8), can be negative even if the forecasts are better than climatology -- that is,

even if $\frac{\pi_{12}}{\pi_{22}} < \frac{P_1}{P_2}$; for unless weather forecasts are perfect, $1/\pi_2\pi_{22}$ can be shown to be greater than $\frac{1}{P_2}$.[1] Thus if the decision maker goes out only when good weather is forecast he will not suffer losses as often, but more total days will elapse (days he goes out plus days he stays in) before he completes the job. If C_L is large and C_F is small, forecasts need be only a little better than climatology to be useful. But if C_L is small relative to C_F , forecasts must be quite good before they have value.

It should be stressed here that even if the calculations yield a negative $V(I_f)$, the forecasts do not necessarily have a negative value. If the decision maker behaves rationally, he will ignore the forecasts, use climatology, and $V(I_f)$ will be zero, not negative. However, if he insists on using the forecasts he will do worse than he would if he ignored them, and $V(I_f)$ really will have negative value for him.

As with the general-purpose framework, it is straightforward to put an upper bound on the value of improvements in weather forecasts. Assume that weather forecasts are perfect. Thus $\pi_{11} = \pi_{22} = 1$, $\pi_1 = P_1$, and $\pi_2 = P_2$. Call this perfect weather information I_∞. From common sense as well as Equation (4) it is clear that if forecasts are perfect, then $\bar{N}_L = 0$. And since it is certain the decision maker will go out and finish the job on the first good day, $\bar{N}_T$ must equal the expected number of days to the first good day. Both Equations (5) and (3) give $\bar{N}_T = 1/P_2$. Thus for perfect forecasts $C(I_\infty) = C_D + C_F \frac{1}{P_2}$. The value of improvements in forecasts cannot exceed $C(I_f) - C(I_\infty)$.

[1] From the preceding footnote we know that $P_2 > \pi_2\pi_{22}$, if bad weather is forecast at all.

In the preceding calculations, the job could be completed in one day of good weather. Clearly, if a job requires M days in good weather, and does not require that the work be done on adjacent days, N_T and N_L in the equations above must be multiplied by M. If direct costs are calculated on a daily basis, they also must be multiplied by M. The analysis will be made clear in the case study which follows.

AN EMPIRICAL EXAMPLE: ROOFING CONTRACTING

Most contractors will not undertake roofing operations if there is high wind or more than light rain. Winds over 20 mph often blow materials away and slow down progress. Rain can damage the underlayer of felt on the sheeting and cause warping later on. Either wind or rain can make work hazardous for the laborers.

If a contractor does not expect wind or rain he arranges for his men to report to work at the job site; if the day turns rainy or very windy, he usually has them go home. Ordinarily, the contractor is able to cancel his crew with no payment if the work has not begun; but once it has, he is obligated to pay his men for at least a quarter of a day. If work is canceled after a quarter of a day, he must pay for the fraction of the day actually worked. Thus it is extremely difficult to calculate the average costs to the contractor of cancelling work. If he cancels before work is started, it costs him nothing but possible ill will. If he cancels after work is started, he does receive some work for his money, although if the rain is sudden and heavy some of the work may have to be redone later. In our analysis we will ignore any work which may be done on days when work is cancelled, and any materials cost on those days. We shall

assume that on 80 per cent of the days when work is cancelled, work starts beforehand, and the costs are equal to a quarter of the day's wages.

A typical roofing project involving work on a number of structures may take about four days of good weather for completion ($M = 4$), a wage bill of about $1,000 a day, and a total materials cost of about $6,000. Thus:

$$C_L = \frac{4}{5} \times \frac{1}{4} \times 1{,}000 = 200 \quad .$$

Overhead for a roofing company may run about $150 a day:

$$C_F = 150 \quad .$$

Figuring total direct costs for the completed job, we have:

$$C_D = 4 \times 1{,}000 + 6{,}000 = 10{,}000 \quad .$$ [1]

If we ignore the problem of wind and consider only rain in our analysis, we find that the relevant probability in parts of Southern California in early spring is about .2. If the contractor has no meteorological information but climatology and thus goes out every day, cancelling work on days that turn out rainy, it will take an average of five days to complete the job:

$$\bar{N}_T = M \frac{1}{P_2} = 4 \times \frac{1}{.8} = 5 \quad .$$

On one of these five days work will be cancelled:

$$\bar{N}_L = M \frac{P_1}{P_2} = 4 \times \frac{.2}{.8} = 1 \quad .$$

---

[1] $M = 4$ is the required number of days of work to complete the job.

And the average cost per completed job will be:

$$C(I_o) = C_D + \bar{N}_T C_F + \bar{N}_L C_L = 10{,}000 + 5 \times 150 + 1 \times 200 = 10{,}950 \quad .$$

Because direct costs comprise a large proportion of total costs and will not be incurred on days when work is not attempted or cancelled, forecasts must be very good to make any marked reduction in roofing costs. The forecast structure we shall consider is characterized by probabilities:

$$\pi_{22} = .96, \quad \pi_{12} = .04, \quad \pi_2 = .7 \quad .$$

If the contractor bases his decisions on these forecasts:

$$\bar{N}_T = M \frac{1}{\pi_2 \pi_{22}} = 6, \quad \bar{N}_L = M \frac{\pi_{12}}{\pi_{22}} = \frac{1}{6} \quad .$$

Thus an average of six days would be required per completed roofing job. Over a span of six jobs (36 days) there would be, on the average, one day when work is cancelled and 11 when no work was attempted. The average cost for each completed project would be:

$$C(I_f) = 10{,}000 + 6 \times 150 \times \frac{1}{6} \times 200 = 10{,}933 \quad .$$

The saving due to weather forecasts would be about $17 a project. Indeed, even if weather forecasts were perfect, the cost per project would be:

$$C(I_\infty) = 10{,}000 + 5 \times 150 = 10{,}750 \quad .$$

Thus weather forecasts, even very good ones, are worth little to roofing contractors because although the weather definitely affects roofing work, there is not much gain in being able to predict bad weather. Only small losses are incurred when bad weather is not predicted.

A POSTPONEMENT MODEL WITH A DEADLINE

We now modify the analysis of the preceding Section by assuming that, instead of having an indefinite period of time in which to complete the job, the decision maker must complete it by some time T or else incur a penalty C_P. The penalty can take several forms. It might be a full or partial refund of any gross revenue paid the decision maker who has undertaken the contract but is not now required to complete it. This case would arise if completion of the job after time T had no value to the person letting the contract. It might be the expected cost of eventually completing the job, with no deadline but with some higher cost per day C_F applying after time T. Or, there might be a series of dates at which one or more of the costs change, in which case the first penalty C_P simply represents the expected cost of fulfilling the contract when the job has not been completed by time T, the calculation of this quantity taking into account all the possible cost changes which might occur. The formal analysis of the deadline problem is the same in every case. We could link together a series of analyses of individual deadline problems to reflect a wide variety of possible patterns of change in the various costs over time.

Consider a job which has not been completed and k days remain before the deadline. In deciding what action to take on that day, the decision maker must take into account the expected remaining cost of fulfilling the contract if he should fail to complete the job on that day -- which reflects, of course, the chance that the job will never be done and the penalty will have to be paid. Let this expected cost be $D(k - 1)$. It is assumed, as before, that there is no persistence in weather states. If forecasts are

available, D(k - 1) refers to the expected cost <u>before</u> the forecast is received for day k-1 (the (k-1)th day from the deadline). On these assumptions, the decision maker faces the following cost matrix on day k:

	W_1(rainy)	W_2(fair)
a_1(nogo)	$C_F + D(k-1)$	$C_F + D(k-1)$
a_2(go)	$C_F + C_L + D(k-1)$	$C_F + C_D$

If we knew the value of D(k-1), the matrix above would present a decision problem which could be handled by applying the theory developed for single-stage decisions in Section I. The complication arises because we do not know the value of D(k-1) unless we have already solved the problem facing the decision maker when he has k-1 days left, this cannot be solved without knowing the answer for k-2 days, and so on. But this sort of difficulty points the way to its own resolution: we can solve the problem for the case when there is one day left, and work backwards. For any particular problem in which numerical values for the various costs and probabilities are available, we can do this by simply carrying out the computations. This statement provides little insight into the qualitative nature of the solution, however -- how it depends on k, the probabilities, and the various costs. The analysis of the following problem should provide this insight.

We begin by considering how the optimal action for the k^{th} day from the deadline depends on D(k-1) and on the forecast received -- or, to put it another way, how the optimal <u>rule</u> of action depends on D(k-1). If forecast f_2 (fair) is received, the expected cost if action a_1 is taken is

$$C_{12} = C_F + D(k-1) \quad ,$$

and for action a_2:

$$C_{22} = C_F + \pi_{22} C_D + \pi_{12}\left[C_L + D(k-1)\right] \quad .$$

It will pay to take action a_2 on forecast f_2 if

$$C_{22} < C_{12} \quad ,$$

which requires

$$D(k-1) > C_D + \frac{\pi_{12}}{\pi_{22}} C_L \quad .$$

Otherwise, action a_1 should be taken on forecast f_2. If forecast f_1 is received, it will pay to take action a_2 only if $C_{21} < C_{11}$, which will be found to require

$$D(k-1) > C_D + \frac{\pi_{11}}{\pi_{21}} C_L \quad .$$

If the forecasts are any good, $\frac{\pi_{12}}{\pi_{22}}$ is less than $\frac{\pi_{11}}{\pi_{21}}$. Therefore, if $D(k-1) > C_D + \frac{\pi_{11}}{\pi_{21}} C_L$, it pays to go out regardless of the forecast; if $D(k-1)$ is less than $C_D + \frac{\pi_{11}}{\pi_{21}} C_L$, but greater than $C_D + \frac{\pi_{12}}{\pi_{22}} C_L$, it pays to go out only on a favorable forecast; while if $D(k-1)$ is less than $C_D + \frac{\pi_{12}}{\pi_{22}}$, it does not pay to go out even if the forecast is favorable. Call these rules of action α_1, α_2 and α_3 respectively, and tabulate them thus:

	f_1	f_2
α_1	a_2	a_2
α_2	a_1	a_2
α_3	a_1	a_1

We will refer to the set of values of D(k-1) for which α_1 is the optimal rule of action as region I, and define regions II and III analogously.

Setting aside temporarily the results just developed, we now consider in turn the cost implications of following each of the rules of action throughout, regardless of the region in which D(k-1) happens to fall. First, suppose rule α_1 is followed throughout. Then the expected cost on day k is related to the expected cost on day k-1, by

$$D_I(k) = C_F + P_2\, C_D + P_1\left[\, C_L + D_I\,(k-1)\right] \quad , \tag{9}$$

for the fixed cost C_F will be incurred regardless of the weather on day k; C_D will be incurred if the weather is fair and the job is completed; and C_L will be incurred if the weather is rainy and the attempt to do the job fails. If the job is not completed by time T, the expected cost $D_I(0)$ of fulfilling the contract will be C_P, since paying the penalty will be the only way of fulfilling the contract. This provides us with the initial condition for the first-order difference equation in $D_I(k)$. The solution of Equation (9) is:

$$D_I(k) = (C_D + \frac{1}{P_2}\; C_F + \frac{P_1}{P_2}\; C_L) + P_1{}^K\left[C_P - (C_D + \frac{1}{P_2}\; C_F + \frac{P_1}{P_2}\, C_L)\right] . \tag{10}$$

As k gets very large, P_1^k gets very close to zero, and $D_I(k)$ is approximately equal to the first term in this expression, which we will denote by $D_I(\infty)$. This will be recognized from the analysis of the preceding section as the expected cost of completing the job, following rule α_1, when there is no deadline. It is entirely reasonable that facing a deadline indefinitely far in the future is the same as facing no deadline at all. Note also that the expression in brackets is the difference between the penalty cost C_P and the no-deadline cost $D_I(\infty)$. If the former is the larger of the two, the expected cost grows as the deadline nears (k gets smaller). Otherwise, the expected cost declines as the deadline nears. Of course, the <u>total</u> cost of the contract grows as time passes and the job is not done; but the expected <u>remaining</u> cost (which is what is relevant to the decision to attempt or not attempt the job) will decline as the deadline nears if C_P is less than $D_I(\infty)$.

Now suppose that rule α_2 is followed throughout. Then $D_{II}(k)$ is given by

$$D_{II}(k) = C_F + \pi_2 \pi_{22} C_D + \pi_2 \pi_{12} C_L + (\pi_1 + \pi_2\pi_{12}) D_{II}(k-1) \quad , \qquad (11)$$

since C_F will be incurred regardless of the forecast; C_D will be incurred if the forecast is favorable and the weather good; C_L will be incurred and the job will not be done if the forecast is for fair but the weather is rainy; and, regardless of the weather, the job will not be done if the forecast is for rain. The solution to the difference equation for $D_{II}(k)$ is

$$D_{II}(k) = (C_D + \frac{1}{\pi_2\pi_{22}} C_F + \frac{\pi_{12}}{\pi_{22}} C_L) + (\pi_1 + \pi_2\pi_{12})^k \Big[C_P - (C_D + \frac{1}{\pi_2\pi_{22}} C_F + \frac{\pi_{12}}{\pi_{22}} C_L)\Big] \quad . \qquad (12)$$

The first term is the no-deadline cost $D_{II}(\infty)$, which applies if rule α_2 is followed continuously in the no-deadline case. Whether this quantity is greater or less than C_P determines whether $D_{II}(k)$ rises or falls as the deadline approaches (as was true with $D_I(k)$).

If rule α_3 is followed throughout, the expected cost on day k is clearly

$$D_{III}(k) = C_F + D_{III}(k - 1) \quad , \tag{13}$$

and the solution to the equation is

$$D_{III}(k) = C_P + k\, C_F \quad . \tag{14}$$

Of course, α_3 can never be the optimal rule of action in the no-deadline case, since it involves incurring the cost C_F indefinitely.

We can now combine the difference equations just developed with the determination of the regions in which each rule of action is optimal, in order to generate D(k) for all k. Suppose, for example, C_P is in region I, but $D_I(\infty)$ is in region II. Then the equation for $D_I(k)$ can be used to generate D(k) until for some k = k', D(k') lies in region II for the first time. At this point we switch to $D_{II}(k)$ for generating D(k), in the following manner:

$$D(k) = D_{II}(\infty) + (\pi_1 + \pi_2\, \pi_{12})^{k-k'} \left[D(k') - D_{II}(\infty)\right]; \text{ for } k \geq k' \quad . \tag{15}$$

That is, we regard D(k') as a penalty cost incurred if, following rule α_2, the job is not completed when k' days remain to complete the job. As long as this new equation, associated with rule α_2, generates values of D(k)

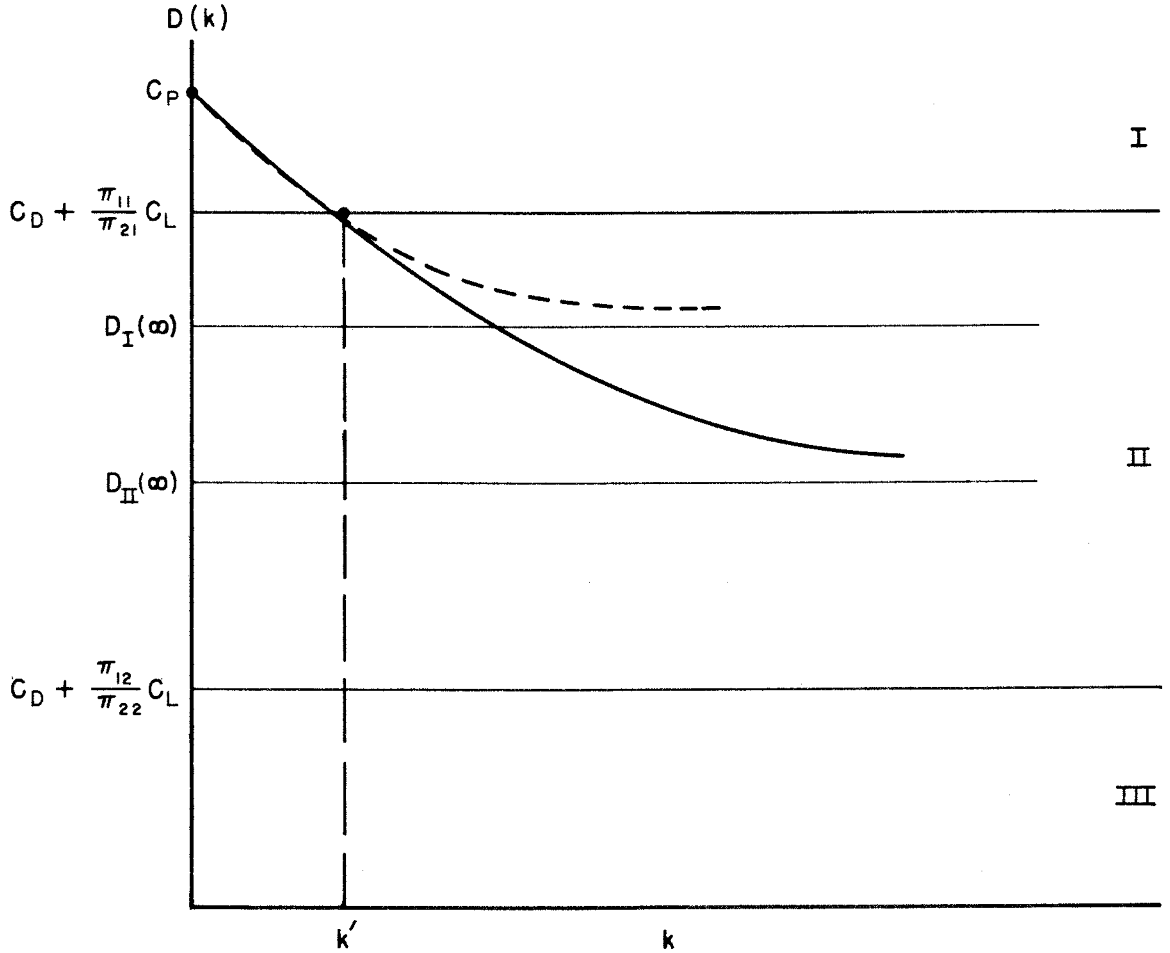
D(k)
C_P
$C_D + \frac{\pi_{11}}{\pi_{21}} C_L$
$D_I(\infty)$
$D_{II}(\infty)$
$C_D + \frac{\pi_{12}}{\pi_{22}} C_L$
I
II
III
k′
k

Fig. 2

which fall in region II, we continue to apply it. If another rule change were required for some large value of k, we could employ the same trick again. Actually, in the assumed case, there will be no further rule change as k enlarges. The path of D(k) is shown in Figure 2.

It is easy to see the reason there will be no further rule change in this case. Whenever a no-deadline cost lies in its own region (e.g., $D_{II}(\infty)$ lies in region II), it must be below the other no-deadline cost. For, in the absence of a deadline, the no-deadline cost is the expected cost of completing the job for all k. If $D_{II}(\infty)$ lies in region II, then rule α_2 is the rule to follow throughout in the no-deadline case. Therefore, $D_I(\infty)$ cannot possibly lie below it, for this would imply that following rule α_1 throughout would be cheaper. Finally, it is easy to see that $D_{II}(\infty)$ is greater than $C_D + \frac{\pi_{12}}{\pi_{22}} C_L$, which implies that neither $D_{II}(\infty)$ nor $D_I(\infty)$ can ever lie in region III.[1]

[1] These statements can be proved mathematically as follows: make the substitution

$$C_F = \pi_2\pi_{22} C_L \left(\frac{\pi_{11}}{\pi_{21}} - \frac{\pi_{12}}{\pi_{22}} + Y\right)$$

in the expressions for $D_I(\infty)$ and $D_{II}(\infty)$; and in the expression for $D_I(\infty)$, substitute $\pi_1\pi_{11} + \pi_2\pi_{12}$ for P_1 and $\pi_1\pi_{21} + \pi_2\pi_{22}$ for P_2. Then with a little manipulation, one reaches

$$D_I(\infty) - C_D - \frac{\pi_{11}}{\pi_{21}} C_L = \frac{\pi_2\,\pi_{22}}{\pi_1\pi_{21} + \pi_2\pi_{22}} C_L Y \quad ; \quad \text{and}$$

$$D_{II}(\infty) - C_D - \frac{\pi_{11}}{\pi_{21}} C_L = C_L Y \quad .$$

This shows that if Y is greater than zero, both quantities are in region I, and since $\frac{\pi_2\pi_{22}}{\pi_1\pi_{21} + \pi_2\pi_{22}}$ is less than one, $D_I(\infty)$ lies below $D_{II}(\infty)$; while

The qualitative behavior of the path of $D(k)$ -- whether it rises or falls, and the regions through which it passes -- can be determined by observing in which regions $D_I(\infty)$, $D_{II}(\infty)$ and C_P lie. For large k, $D(k)$ will be close to the lower of $D_I(\infty)$ and $D_{II}(\infty)$; for small k, it approaches C_P and the path will be monotonic. Of course, whenever the path of $D(k)$ crosses a boundary between regions, the equation governing it changes. Some possible paths are shown in Figs. 3 through 6. Figures 5 and 6 illustrate cases which may seem paradoxical. The decision maker apparently gives up on doing the job before the deadline arrives, in spite of the fact that he was attempting to do the job when the deadline was far in the future! Although the person who wants the job done may have made a mistake by imposing such a low penalty cost for failure, there is nothing nonsensical about this kind of a pattern. The prospect of incurring a long series of C_F costs is sufficient to give the decision maker an incentive to attempt the job when the deadline is far in the future. But if he has bad luck and fails to complete the job early, he eventually reaches a situation where the additional cost he expects to incur if he fails to complete the job on a particular day is relatively low -- and this implies a diminished incentive to attempt the job. The expected cost of completing the contract with no forecasts available is given, for all k, by the equation for $D_I(k)$, with C_P as the initial condition $D_I(0)$. By comparing the value

---

if Y is negative, both lie in region II and $D_{II}(\infty)$ is the lower of the two. Finally, $D_{II}(\infty)$ exceeds $C_D + \frac{\pi_{12}}{\pi_{22}} C_L$ by $\frac{1}{\pi_2 \pi_{22}} C_F$, and this, together with the first result, proves the last statement.

given by this equation with the value of $D(k)$ we can determine, for any particular number k of days allowed before the deadline, the value of the forecasts. For large k, the value of the forecasts is approximately the same as in the no-deadline case. For small k, things may be different. The forecasts may have value even when they have none in the no-deadline case, or they may fail to have value whereas they do have it in the no-deadline case.

Graphs of the type shown in Figs. 1 through 6 also make it a simple matter to determine the effect of changes in the deadline date on the expected cost of completing the contract.

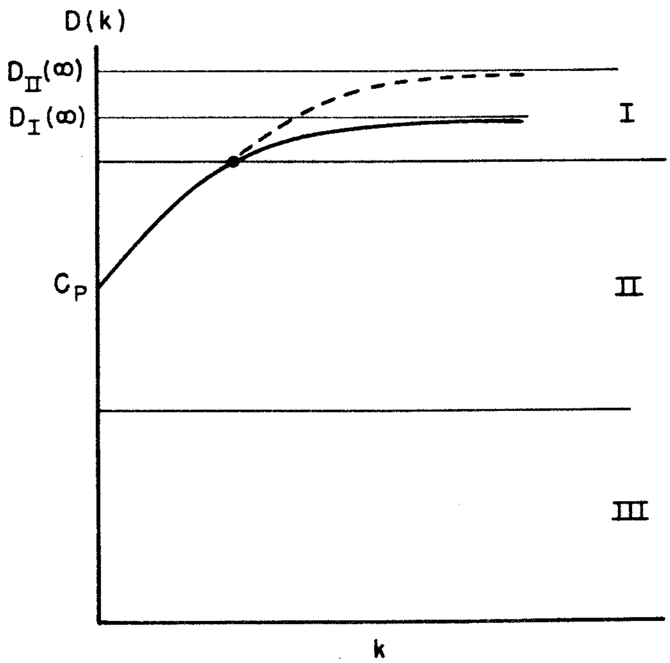

Fig. 3

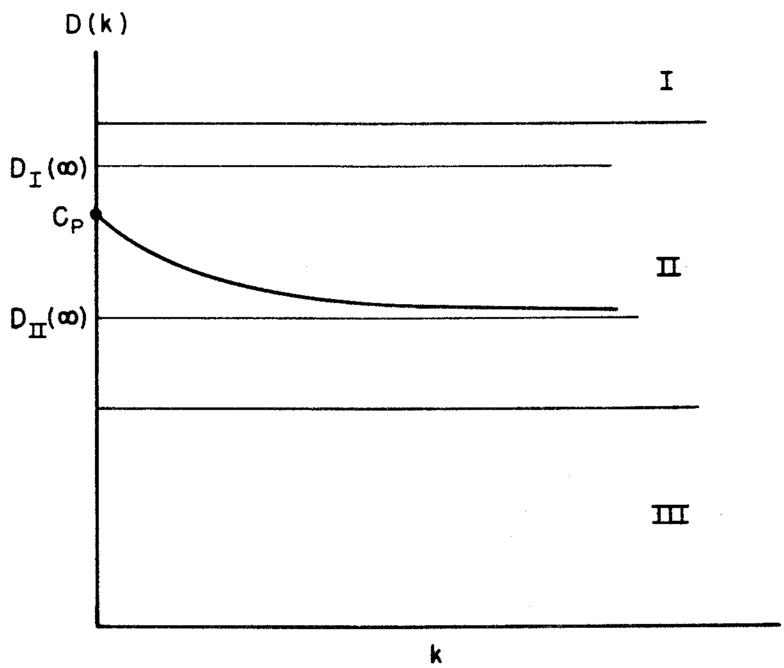

Fig. 4

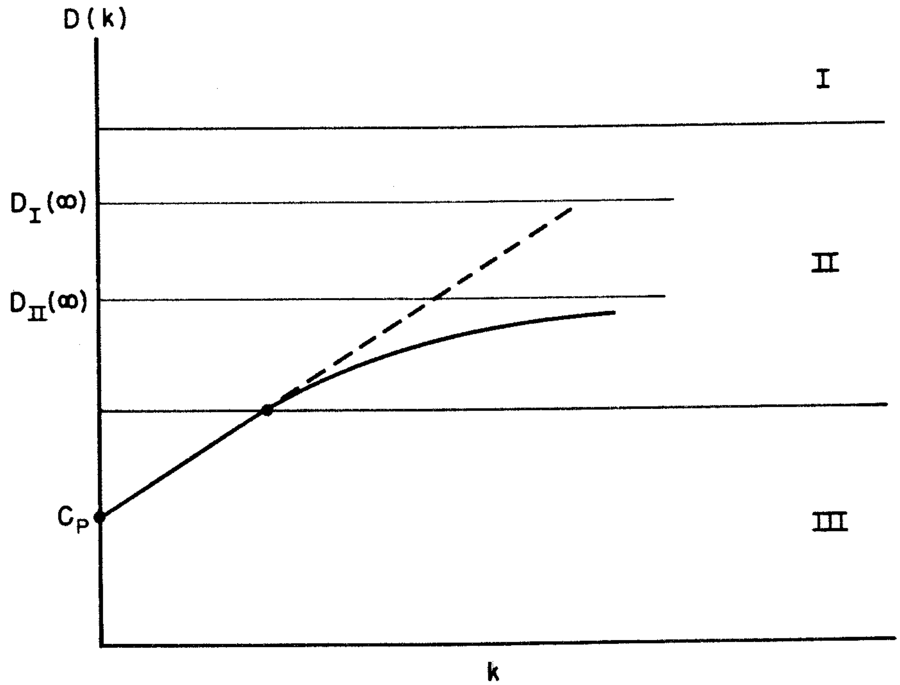

Fig. 5

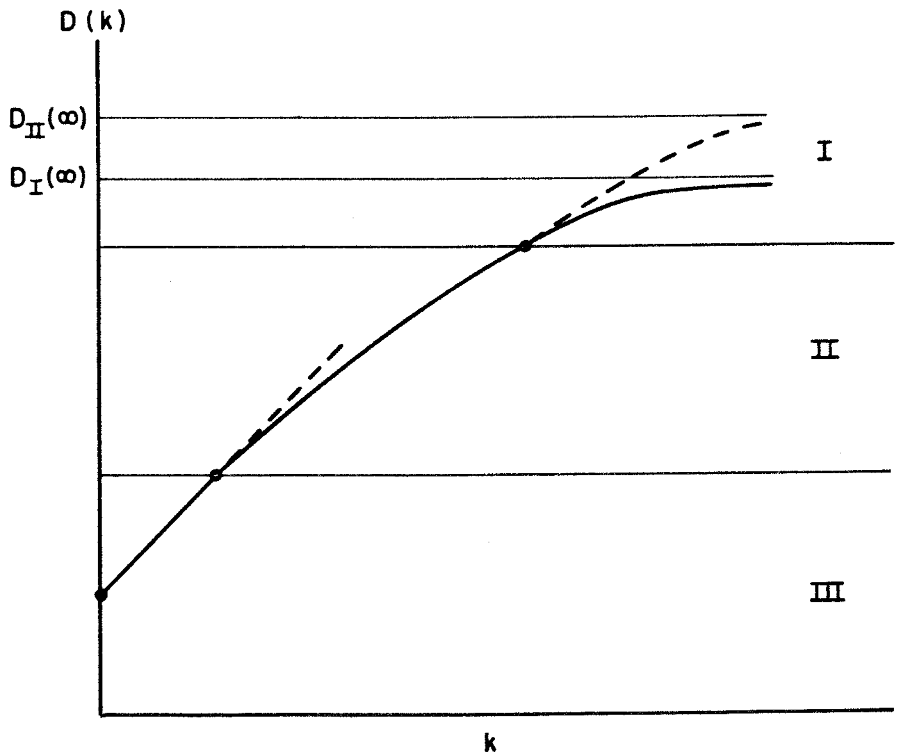

Fig. 6

IV. THE MAKING OF FORECASTS: OPTIMAL FORECAST STRUCTURES

INTRODUCTION

In the preceding sections, we examined in some detail a few situations in which weather information had value. Disregarding the purely intellectual interest in knowing what the future holds, we stressed the point that weather forecasts have value because they enable people to make better decisions. In making this point we focused our attention on the decision problem facing the consumer of weather forecasts, taking as given the joint distribution of weather states and forecasts summarized by the information vector I_f. For any particular I_f, we showed how optimal rules of action can be developed and the value of the forecasts determined. In the present section, we extend the analysis to the question of how I_f is determined. We note that the forecaster has to decide upon a forecast structure, which has been defined in Section I as a function relating his forecasts to observations on meteorological variables, w and z, the time to which the forecasts apply, t, and the location to which they apply, L, $f = \eta(w, z, t, L)$. Naturally, the problem of choosing a forecast structure is largely a problem in meteorology. The forecaster's knowledge of the relations among different meteorological variables and of the way in which weather changes in space and time determines the set of alternative η's which are available to him. The range of possible choices is narrowed considerably when the particular observations which are available and the analytical procedures which the forecaster employs are specified. But the specification of these things merely determines the amount and type of (uncertain) knowledge of future weather which is available to the

forecaster. The question of how this knowledge should be conveyed to the consumers of the forecast is examined in this section. Though obviously only a small part of how to design an optimal forecast structure, it is that part of the problem where the relevance of economic considerations to the choices made can be shown most clearly.

A DIGRESSION ON THE CONCEPT OF FORECASTING

In subsequent parts of this section, we will present the economic criteria for an optimal forecast structure and illustrate the application of those criteria in particular cases. First, however, we will digress briefly to discuss the concepts of forecasts and forecasting. It will be noted that the concept of a forecast employed in the present section differs from that employed elsewhere in this paper. The relationship between our formal characterization of the process by which forecasts are made and the ways they are actually made in the real world will be discussed.

A forecast has been defined (p. 3) as a signal with which the decision maker can associate known objective probabilities of the future weather states which are relevant to his decision. We have not attached significance to the exact words or numbers used to convey the signal to the decision maker, and in cases where we have applied a label to a forecast ("dry" or "rainy"), the label has not itself contained the information upon which the decision maker is assumed to act -- that is, it has not itself contained the relevant probabilities. We have assumed that the

decision maker knows the true relation between the signals he receives and the probabilities of the weather states, presumably through long experience with the forecasts. This assumption is very convenient, perhaps essential, in an analysis of the value of forecasts in which the set of forecasts and the probabilities with each forecast are taken as given. But the question of how the labeling of the forecasts may affect their value cannot be pushed aside in a discussion of how the forecaster's knowledge of the weather should be conveyed to the decision maker. Intelligibility is obviously an important consideration. Or, to put it another way, we must not ignore the costs of information processing which the labeling of the forecast imposes upon the decision maker, or the possibility that the decision maker may be unaware of the relevance of the forecasts to his problem if the labeling is very inappropriate.

If we adhere strictly to the idea that the decision maker, through long experience with the signals he receives, can attach the correct probabilities to them himself, we must concede that many signals may satisfy our definition of forecasts which do not carry a label suggestive of the signals' relevance to future weather states. For example, a report that a hurricane is over the ocean says nothing explicit about the possibility that it will strike a particular section of the coast -- but a decision maker who has lived for a long time in that area may be able to assign probabilities to the report just as easily as he could to an explicit prediction that the hurricane will pass near him. A trucker, told that there is a large storm along a certain route, may be able to attach probabilities to the storm's still being there a few hours hence as easily as he could to an explicit prediction that it will still be there. A decision maker skilled in the

reading of weather maps may find it as easy to estimate the probabilities of the weather states relevant to his decision by looking at the map as by reading the forecast.[1] In deriving from his experience the probabilities to assign to various signals which are not phrased as predictions, the decision maker is not doing anything essentially different from what we have assumed him to be doing when he receives, for example, a categorical forecast that it will rain. In each case the true significance of the signal -- the probabilities of future weather states associated with it -- is not conveyed by the words in which it is stated (except in the case where the categorical forecast is perfect). The distinction between the cases is one of degree. In each case the decision maker himself has to exert some effort to convert the signal to a usable form; the question is just <u>how much</u> effort he has to exert and whether it is desirable for him to exert this effort instead of having the forecaster exert it for him. For our purposes here, we define a forecast as any signal provided to the decision maker which is based on something more than climatology -- and leave open the question of whether the decision maker understands the signal correctly.[2]

---

[1] In all of these cases, we do not mean to suggest that the probabilities assigned to these various signals would necessarily be the same as the probabilities assigned to the explicit forecasts. Our point is that the difficulty of assigning the probabilities is not much greater in the one case than in the other.

[2] This discussion of the ability of the decision maker to make use of signals which do not make explicit predictions raises some of the same issues that are discussed in Appendix 2 of Section I, where it is noted that the decision maker may be able to do his own forecasting by exploiting the fact that weather states tend to persist. In both cases, analysis of the situation is complicated by the difficulty of deciding what costs to the decision maker are involved when he does his own forecasting.

We now turn to a possible objection to the characterization of the forecasting process contained in our definition of a forecast structure. That definition indicates that forecasting is a process in which meteorological variables and data on the time and place to which the forecast refers are mapped into a forecast by applying a particular functional form η. It might be objected that this implies that forecasting is a strictly deterministic process which can be reduced to a set of fixed rules, whereas in fact many forecasters operate by methods which they cannot entirely explain. Such forecasters are practicing an art rather than a science, and among those who operate in this way are some with very good records. We are prepared to admit that a forecaster is not a machine, and that a complete analysis of how a particular forecaster happens to come up with a particular forecast on a particular occasion may be impossible. Nevertheless we feel that our characterization of the process is a convenient one. The main point we want to emphasize is the existence of a range of choice which is open to the forecaster when he develops forecasts out of the observations available to him. The range of choice is most apparent when the forecaster's analysis of his data is complete and a decision must be made as to what weather information is to be transmitted to the decision maker, and how. Even a forecaster who relies heavily on intuitive procedures which he cannot fully describe should be able to provide more detail about what is likely to happen in certain weather dimensions and less to others.

We would argue that an important range of choice also exists at earlier stages in the process which generates forecasts. For example, there must be alternative sets of observations available at the same cost and which have differential usefulness for making various predictions. There

must be observations of particular importance in providing tornado warning but making a relatively small contribution toward forecasting the total rainfall that will occur over a period of a week, and vice versa. Perhaps there are similar choices to be made among alternative techniques of analysis. Consideration of all of these aspects of the design of forecast structures would require a discussion of forecasting as a problem in meteorology, and in this context the objections to our simple deterministic view of forecasting would carry more weight. But for the present discussion, the characterization is adequate.

Furthermore, there exist forecasting methods which conform precisely to our concept of a forecast structure. In objective (or "numerical") forecasting methods, observations on a set of weather variables are combined in a definite way to yield a quantitative indicator of future events in some weather dimension. The formula for combining the variables is derived by statistical techniques, while the variables considered and the functional form of the relationship are suggested by theoretical considerations and the availability of the required observations. Forecasting methods of this kind are still in the process of development, but some results have been obtained which are clearly of a quality comparable to that of the more typical subjective forecasting methods. The fact that the dependence of the forecaster's knowledge of future weather upon the observations is made explicit in the objective scheme makes possible a precise description of the set of alternative forecast structures which can be based upon the scheme. Later in this section, we will give examples of the determination of optimal forecast structures based upon an objective scheme. First, we must present the economic criteria employed in such a determination.

CRITERIA FOR OPTIMAL FORECAST STRUCTURES

The first step toward developing the criteria for optimal forecast structures is to set forth the reason why the problem exists. Why should the forecaster convey to decision makers less than his total knowledge of future weather events? The reason is that there are costs associated with the process of conveying the information, and there may also be restraints on the amount of information which can be conveyed. The weather forecast cannot occupy more than a certain amount of space in the newspaper or take up more than a certain amount of radio or TV time. Nor can a public or private forecaster spend an indefinite amount of time talking on the telephone or in person to his consumers. All of these limitations may be thought of as restraints on the "channel capacity" available for transmitting weather information to the decision maker -- restraints which it may be possible to relax, but only at some cost. In addition, the amount of information which can be conveyed is limited by the information-processing capacity of the decision maker, and the costs of increasing that capacity. In simpler words, the forecast must be intelligible to the decision maker without his having to go to great effort to understand it. The greater the quantity of information transmitted to him, the greater the problem he may have in reducing it to a useful form. The forecaster must avoid giving such sophisticated forecasts that the decision maker has to have a degree of meteorology to learn that the forecast says, "It looks like a hurricane is coming."

Much of the forecaster's knowledge of future weather relates to weather variables which have no direct relevance to economic activity. For example, few, if any, activities are directly affected by atmospheric pressure. A

decision maker whose costs depend on whether or not it rains at some time or other has no interest in the atmospheric pressure except to the extent that it serves as an indicator of the probability of rain. If the forecaster has already taken atmospheric pressure into account in developing his rain forecast, there is no need to provide the decision maker with data either on the current pressure or on the expected future pressure. At least there is no need to do so if the signal the forecaster makes to the decision maker is a "sufficient signal"[1] -- one which contains all of the information relevant to the decision maker's problem which the forecaster has available. If the signal is not sufficient, knowledge of the atmospheric pressure or any other variable which enters the forecaster's calculations or is correlated with a variable which enters those calculations may be of use to the decision maker.[2] Thus our first principle of forecast structure is that there is no need to provide the decision maker with more than a sufficient signal. And in view of the costs of transmitting the information and/or the decision maker's costs of interpreting it, there may be good reason to avoid giving more than this.

---

[1] Formally, a sufficient signal is a signal f_s such that

$$\hat{\alpha}(f_s) \equiv \hat{\alpha}(f_s, f) ,$$

where f is any other signal which the forecaster can make. The terminology comes from the statistical concept of a sufficient statistic.

[2] This assertion does not depend on the decision maker's having a knowledge of meteorology comparable to that of the forecaster. No matter how much the forecaster knows about future weather, the decision maker may be able to do better by combining some observation of his own with the forecast. From one point of view this is trivially obvious, but the problem of determining a sufficient signal is not trivial, and it is not necessarily easy for the forecaster to provide forecasts which the decision maker cannot improve upon by making his own observations, or by having some of the forecaster's observations available.

The concept of a sufficient signal is closely related to our distinction between the meteorological variables w and z. The former are weather states which affect the decision maker's costs; the latter are other weather variables. Sufficient signals contain all the knowledge the forecaster has about the probabilities of the various w states; they need not convey anything about the z variables. The forecaster may or may not exploit observations on the current state of the weather in developing his forecast. But the value of his forecast derives from the information it provides about the w that will occur in the future.

When the forecast structure is optimized for a single decision maker, the fact that all he requires is a sufficient signal generally makes possible a dramatic reduction in the amount of information which the forecaster should convey to him. But the costs of transmitting or interpreting the information can make it optimal to provide less information than a sufficient signal would provide. For example, if a decision maker can take any one of three actions, but two of the actions involve him in very similar costs in each weather state, the costs of providing the decision maker with information which permits him to distinguish the situations in which each of the two similar actions is optimal may exceed the cost saving made possible by the extra detail. The examples of the determination of optimal forecast structures given in the next section are really examples only of the determination of the simplest possible set of sufficient signals which can be provided to a single decision maker. We do not treat explicitly the costs of transmitting and interpreting the information, although such costs are implicit in our insistence on having the <u>simplest</u> sufficient signals. We assume that they are not so great as to make it desirable to provide signals

which are not sufficient.

When forecasts are provided for more than one decision maker, as is obviously the case in forecasting as a public service at public expense, it seems clear that a signal which is simultaneously sufficient for all decision makers is a very complex signal indeed. This is so because of the very great variety of decision problems to which the weather is relevant, and the variety of ways in which the effect of the weather is felt. It is easy to think of examples where the character of precipitation (rain, snow, or sleet), the amount of precipitation, the occurrence of precipitation regardless of amount, the extreme values of temperature over a certain period, the average temperature, the extreme wind velocity, the extent of cloud cover, visibility, or humidity might define the weather state affecting the decision maker's costs or payoff -- and this is hardly a beginning of a complete list. In principle, the optimization of the forecast structure in order to meet the needs of all consumers simultaneously should reflect all of the information-processing costs imposed on decision makers, the transmission costs, and the relative values of various levels of detail in the forecasts. Detail which is not relevant to any decision maker should be excluded. Whatever relevant detail is excluded from the forecast (as it generally must be), it should be detail which would contribute less to the value of the forecast than any included detail which places equal strain on the effective limit on the forecast -- say, which takes up equal space in the newspaper, if it is newspaper space that limits the amount of information conveyed. If it is possible to relax the restraints on information conveyed at some cost, this should be done as long as the resulting addition to the value of the forecast exceeds the cost of relaxing the restraint. In optimizing a forecast

structure where the forecast is provided at public expense, it is the social value rather than the total private value of the forecasts that should be maximized. It is easy to construct cases where social value and private value are different, but the determination of the social value of particular forecasts is a problem which is difficult or impossible even in principle, let alone in practice.

The statement of these conditions for an optimal forecast structure makes it clear that nothing close to the optimum is likely to be approached in practice. The forecaster has only a very vague idea as to whom his customers are, what activities they are engaged in, how their decisions are influenced by his forecasts, how their information-processing costs are affected by the way in which the forecast is presented, and how their profits and losses depend on the decision they take and on the actual state of the weather. He cannot estimate the social benefits of alternative forecast structures. The best he can do is to meet some of the most important and obvious needs for weather information. Analysis of the problem of designing optimal forecast structures cannot be expected to point out the high road to perfection. It can, however, be expected to point up situations in which important increases in the economic value of the forecasts can be expected to result from fairly simple changes in the forecast structure. The examples and analysis in the remainder of this section should be viewed in this light.

TWO EXAMPLES OF THE DETERMINATION OF OPTIMAL FORECAST STRUCTURES

In this section, we present two examples of the application of the criteria for optimal forecast structures where the structure is optimized to provide forecasts for a single decision maker. The decision problems

are those of the trucking industry of Section I and the movie industry of Section II. The forecast structures are based upon J. C. Thompson's objective scheme for forecasting rainfall in the Los Angeles area.[1] A brief description of this forecasting method is presented in the appendix to this section. The important feature of this method is that all of the knowledge of future rainfall which it produces is summarized in the value of a single variable, called Y_2 in the Thompson article. Empirical probability distributions of the amount of rainfall are available for the different values which Y_2 may take on.[2] The problem of forecast structure design which remains to be solved is how, given the value of Y_2, the distribution associated with that value should be summarized in a forecast. We show how

---

[1] J. C. Thompson, "A Numerical Method for Forecasting Rainfall in the Los Angeles Area," Monthly Weather Review, July 1950, pp. 113-124.

[2] One would need a great deal of experience with the forecasting scheme under stable climatological conditions to have a very accurate estimate of the probability distribution. That is, one would like to have a sample so large that the difference between the sample distribution and the true distribution could be assumed to be very small. In what follows, we employ data from a sample which is not that large, but we ignore the complication introduced by the fact that we are dealing with a sample distribution. The curves which we fit to the data reveal our belief that there is sampling variation in the data, but we treat the curves as if they were the true curves associated with the true distribution, and not estimates. One could argue that this is legitimate for the purposes of analysis of the decision problems of the forecaster and the consumer of the forecasts at a given time, even though as time goes on the estimate of the distribution will change and our calculations of the value of information may not be borne out by historical experience.

There is an additional complication in that the data which we employ are the data from which the forecasting scheme was derived, and hence there is some tendency for the forecasting scheme to give better results with these data than it would with independent data. However, the scheme was tested against independent data, and its performance (as measured by percentage correct, or the "skill score") was very close to that displayed on the original data. We use the original data because the complete distribution for Y_2 and the associated rainfall for the independent data was not presented in Thompson's article.

the solution to this problem depends on the character of the problem facing the decision maker.

Assume first that the forecasts are made for the benefit of the decision maker in the trucking example, who would take action a_1 (protect) if he knew for sure that it would rain more than .15 inches in the period covered by the forecast, action a_2 (do not protect) if he knew for sure that rainfall would be .15 inches or less, and who has no other actions available. In this case, the only thing the decision maker is interested in knowing is whether or not it will rain more than .15 inches. Knowing only this about the decision maker's needs is enough to provide the forecaster with the information he needs to design an optimum forecasting scheme. He knows, for example, that there is no point in providing information as to whether rainfall will be half an inch or two inches, since the decision maker will take the same action in either case. An obvious step toward deciding how to make forecasts for this particular decision maker is to consider the empirical probability (relative frequency) of rain in excess of .15 inches as a function of Y_2. This information for a period of 358 days, is presented in Table 1. It is displayed in graphical form in Fig. 7, where the probabilities are plotted at the mid-points of the Y_2 intervals given in Table 1, and a freehand curve has been drawn through the scatter.

If the forecaster can make a probability forecast -- that is, if he has the means to convey the value of the probability of rain to the decision maker, and the decision maker has the means to understand this signal -- then Fig. 7 solves his problem. He simply determines the value of Y_2, reads off the probability of .16 inches or more of rain from Fig. 7, and conveys this figure to the decision maker. In so doing, he provides the

Table 1*

Y_2 Interval**	Number of days	Days of rain (rain $\geq$ .16 in.)	Relative frequency
Less than 0	54	0	.00
0 to 10	55	0	.00
10 to 20	60	0	.00
20 to 30	45	1	.02
30 to 40	54	3	.06
40 to 50	36	3	.08
50 to 60	21	6	.29
60 to 70	14	4	.36
70 to 80	9	6	.67
Over 80	10	8	.80
Total	358	32	.09

* This table was constructed by counting points on Fig. 5, p. 118, of Thompson's article (*op. cit.*).

** Intervals are inclusive of their upper bounds, exclusive of their lower bounds.

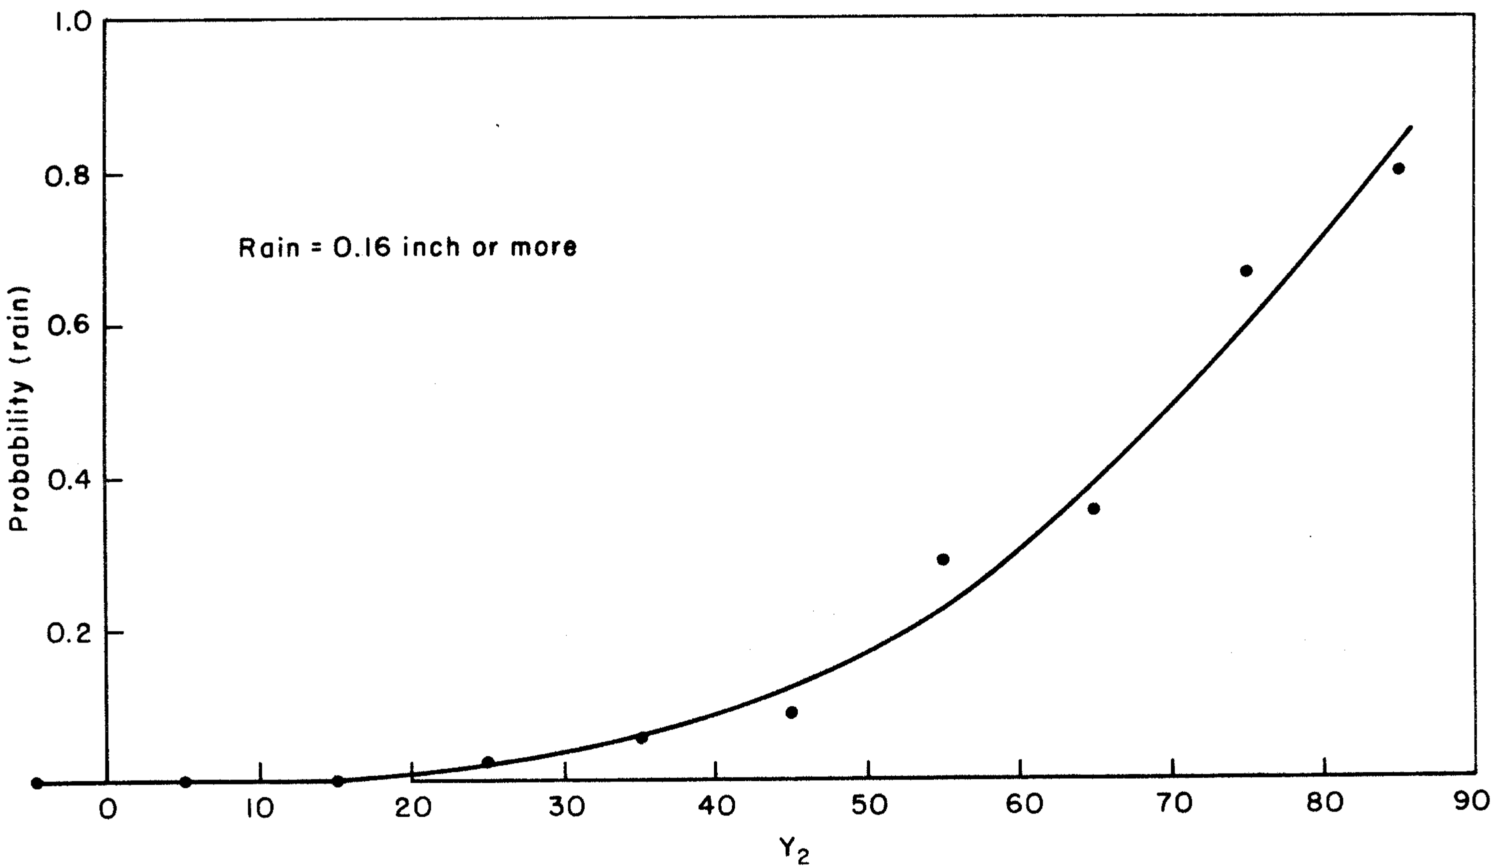
1.0
0.8
0.6
0.4
0.2
Probability (rain)
Rain = 0.16 inch or more
0
10
20
30
40
50
60
70
80
90
Y_2

Fig. 7

decision maker with a sufficient signal -- all of the information relevant to the decision to be made which the forecasting scheme makes available. But suppose that, for some reason, the forecasts must provide the decision maker with the simple statement that there will or will not be rain of the relevant amount during the forecast period. The obvious forecast structure is to choose a cutoff value Y_2^*, forecast rain if the observed value of Y_2 is greater than Y_2^*, and forecast no rain if it is not. The higher the value of Y_2^* chosen, the more frequently the no-rain forecast will be made, but the less accurate it will be when it is made, that is, the more frequently the forecast of no rain will be followed by rain.

Clearly, the forecaster is faced with a variety of possible categorical forecast structures, and he must somehow make a choice among these possibilities. But before considering how he might go about making this choice, let us examine the available alternatives more closely. In Fig. 8, we show the proportion of the time values of Y_2 less than or equal to any Y_2^* occur. Hence, for every value of Y_2^*, the figure shows the proportion of the time that a forecast of no rain would be made. Fig. 9 shows the accuracy of this categorical forecast for every value of Y_2^* -- the proportion of rainy days in days for which Y_2 is less than or equal to Y_2^* -- or, to put it another way, the conditional probability of rain, given that Y_2 is less than or equal to Y_2^*. If we denote the probability that the no-rain forecast is made by π_2, and the probability that rain occurs when the no-rain forecast is made by π_{12}, Figs. 8 and 9 show how these quantities vary as a function of Y_2^*. In Fig. 10, the characteristics of the possible forecasts are presented in a different form. Here we show how the frequency with which the no-rain forecast can be made varies with the

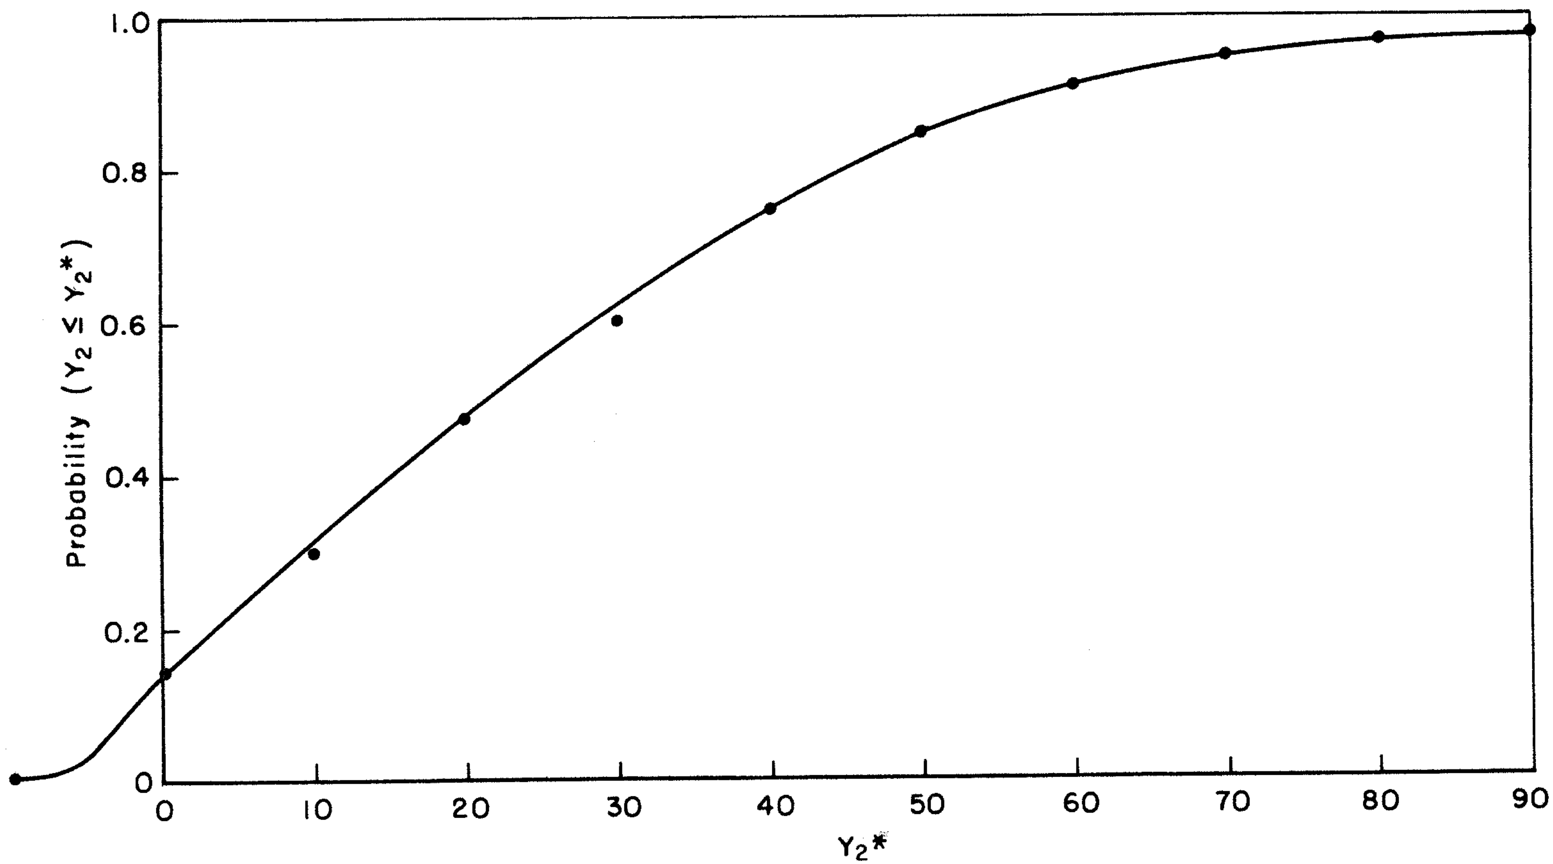
1.0
0.8
0.6
0.4
0.2
0
Probability ($Y_2 \leq Y_2^*$)
0
10
20
30
40
50
60
70
80
90
Y_2^*

Fig. 8

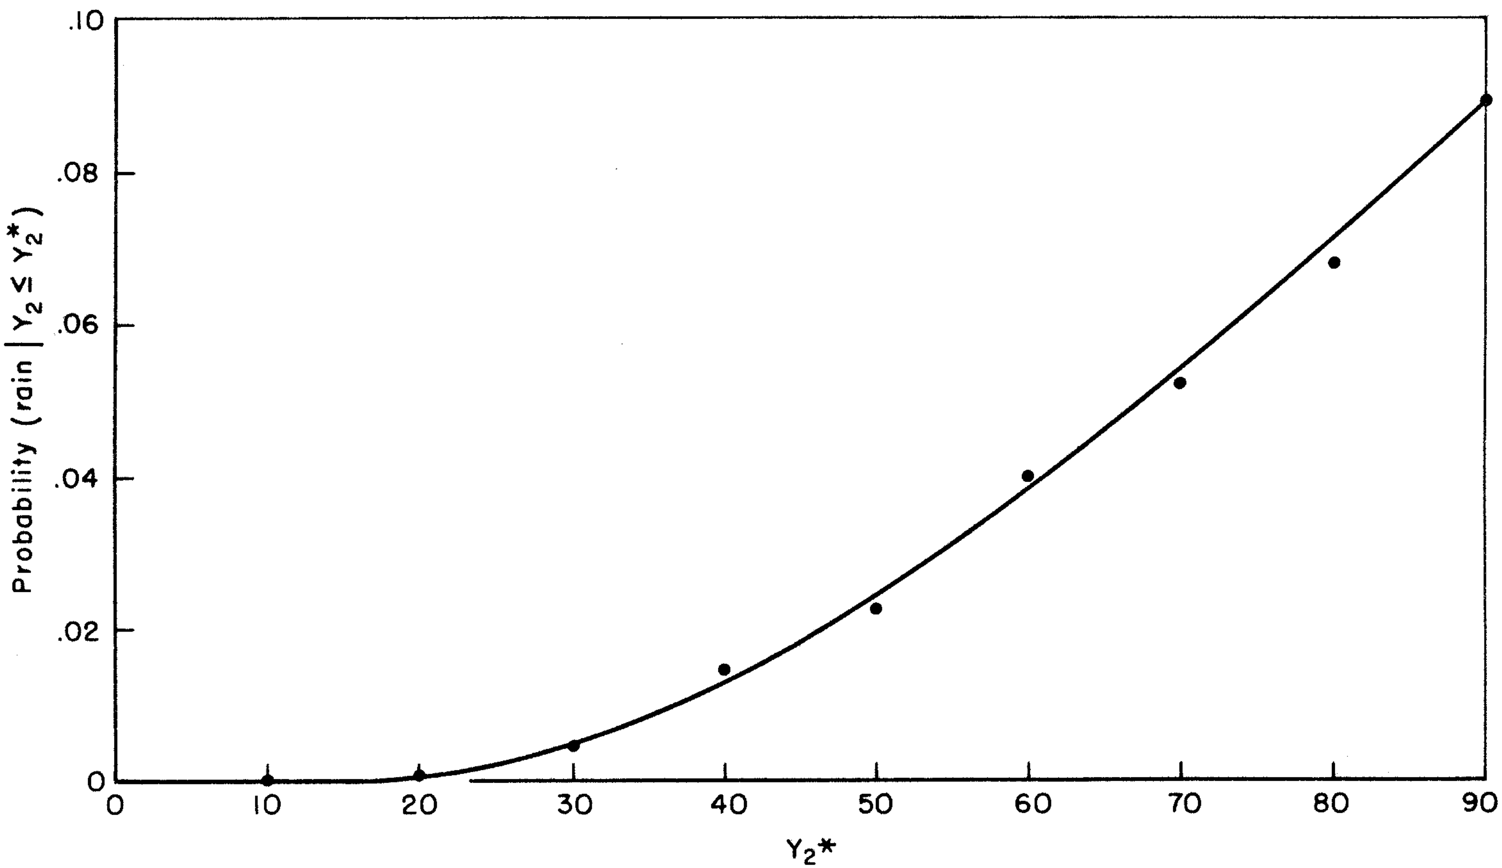
.10
.08
.06
.04
.02
0
Probability (rain | $Y_2 \leq Y_2^*$)
0
10
20
30
40
50
60
70
80
90
Y_2*

Fig. 9

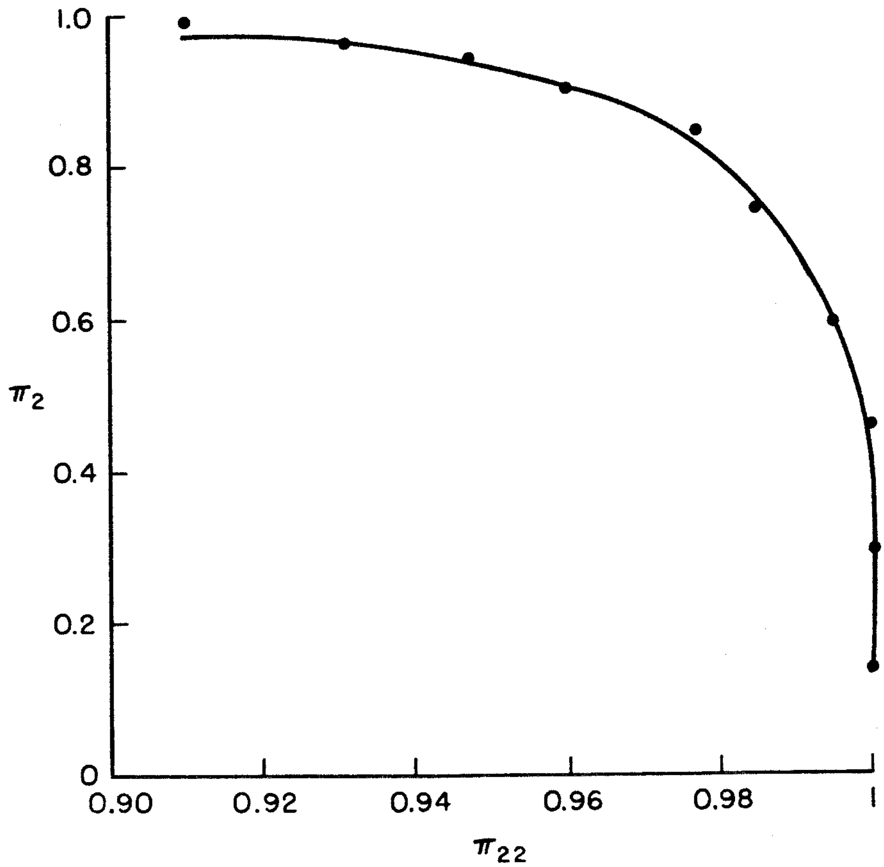

Fig. 10

accuracy required of it. The measure of accuracy on the horizontal axis is now π_{22} , which is simply $1 - \pi_{12}$. Behind the scenes in this relationship between accuracy and frequency is the varying value of Y_2^* : each pair of values (π_2, π_{22}) on the curve corresponds to a particular choice of Y_2^* . Greater and greater accuracy in the forecast of no rain can be achieved by insisting on more and more unequivocal evidence of fair weather, but this unequivocal evidence is available less frequently, and consequently the forecast can be made less often.

It should be noted here that since $\pi_1 + \pi_2 = 1$, π_1 must increase as π_2 decreases when we move along the curve in Fig. 10. As the criterion for a no-rain forecast becomes more stringent, the criterion for a rain forecast becomes less stringent. And thus as π_{22} (the accuracy of the F_2 forecast) increases, π_{11} (the accuracy of the F_1 forecast) must decrease. Our whole discussion could be carried on in terms of the accuracies π_{11} and π_{22} instead of the frequency π_2 and the accuracy π_{22} . That is, we could characterize the alternatives open to the forecaster and the value of the information to the decision maker in terms of these two parameters. This might be more natural as a way of characterizing the quality of a forecast structure, but on the other hand, it seems useful to consider the alternative choices of Y_2^* as corresponding to a set of alternative accuracy-frequency pairs.

To decide among the various ways of making a categorical forecast, the forecaster must have more information on the needs of the decision maker. We have noted above that the decision to be made is of the protect-do not protect type, that is, the cost matrix is of the form shown on p. 8 Section I. Suppose that the decision maker knows the probability of rain

to be P_1 . Then his expected cost if he takes action a_1 (protect) is:

$$C_1 = P_1 C + (1 - P_1) C = C ,$$

and for action a_2 it is:

$$C_2 = P_1 L + (1 - P_1) 0 = P_1 L .$$

He will choose between these actions in such a way as to minimize expected cost; hence he will take a_1 if

$$C_1 < C_2 ,$$

i.e., if

$$C < P_1 L$$

or

$$P_1 > C/L .$$

Otherwise he will take action a_2 .[1] This shows how the decision maker will act, if he is rational, in response to a probability forecast. But notice that he will act in exactly the same way if the forecaster tells him only whether the probability of rain is greater or less than C/L , and his costs will be exactly the same. It follows that this simple distinction provides a sufficient signal. No additional detail about the probability of rain, the value of Y_2 , or the sea level pressure difference between San Francisco and Los Angeles will do this decision maker any good. If there is some cost in conveying this additional information, or if the decision maker may be confused by it, then clearly such information should not be included in the forecast. It appears, therefore, that the best

---

[1] Of course, if $P_1 = C/L$, he is indifferent as to which action he takes.

forecast is a statement as to whether the probability of rain is greater or less than C/L . But the categorical rain - no-rain forecast can serve exactly this purpose if Y_2^* is chosen so that the associated probability of rain (Fig. 7) is C/L . If the decision maker treats the rain forecast as a signal to protect and the no-rain forecast as a signal to dispense with protection, then this categorical forecast does as well as any.

We have now solved the problem of how the objective forecasting scheme should be applied to yield a categorical forecast structure tailored to the needs of a decision maker with a particular cost/loss ratio. It is of some interest, however, to consider how the value of the forecast is affected by departures from the optimum choice of Y_2^* in the particular case of the trucking problem. In examining this question, we will present an alternative method for determining that optimum. In the trucking problem, the climatological probability of rain is greater than the cost/loss ratio, so that in the absence of other information, the decision maker should protect all the time. In this case the value per day of a forecasting scheme which forecasts good weather with frequency π_2 and is in error π_{12} of the times when it makes that forecast has value

$$V = \mathrm{Max}\ [0,\ \pi_2 C - \pi_2 \pi_{12} L]\ .$$

The reasoning behind this formula is very simple. Assuming the forecast has value greater than zero, the decision maker dispenses with protection every time he gets the forecast for fair weather and saves C by doing so; this happens π_2 of the time. Against this saving must be charged the loss resulting from "getting caught" part of the time when the forecast is for good weather, and this happens $\pi_2\pi_{12}$ of the time. Of course, the condition that the forecast have positive value is the same as the condition

that it should pay to dispense with protection when the forecast is for good weather; the probability of rain, given the forecast for good weather, must be less than the cost/loss ratio. An upper bound on the value of information is C, since probabilities must fall between zero and one. This upper bound, however, is greater than the value of perfect information, which is limited by climatology as well.[1] Let us rewrite the formula above as

$$V = \mathrm{Max}[0, \ \pi_2 C - \pi_2(1 - \pi_{22}) L] = \mathrm{Max}(0, \pi_2 C + (\pi_2 \pi_{22} - \pi_2) L] .$$

For every value of V between 0 and C, we can determine the locus of values of π_2 and π_{22} which yield that value of V. Taking $C = 200$ and $L = 5000$, as in the trucking example, we plot some of these loci in Fig. 11. On the same diagram we repeat the locus of attainable points from Fig. 10. The forecaster can choose the method of forecasting which maximizes the value of the information he supplies by finding the point where the curve of attainable forecasts touches the highest possible iso-value of information curve. He can then find the value of Y_2^* corresponding to this pair of probabilities, and act accordingly. Of course, this value of Y_2^* is exactly that derived previously, but we will not introduce the demonstration at this point.

Fig. 11 shows how different forecasts have different values, ranging from 0 to V_{max} which is $104 per day -- each of the forecasts must be regarded as of equal scientific or technical merit, since each represents

---

[1] If $\pi_{12} = 0$, so that rain never occurs when good weather is forecast, then climatology decrees that the best that the forecaster can do is to forecast good weather with the same frequency with which it actually occurs; if he forecasts it more often he is necessarily wrong part of the time. Hence the least upper bound on V is $P_2 C$, where P_2 is the climatological probability of no rain. This is the value of perfect information.

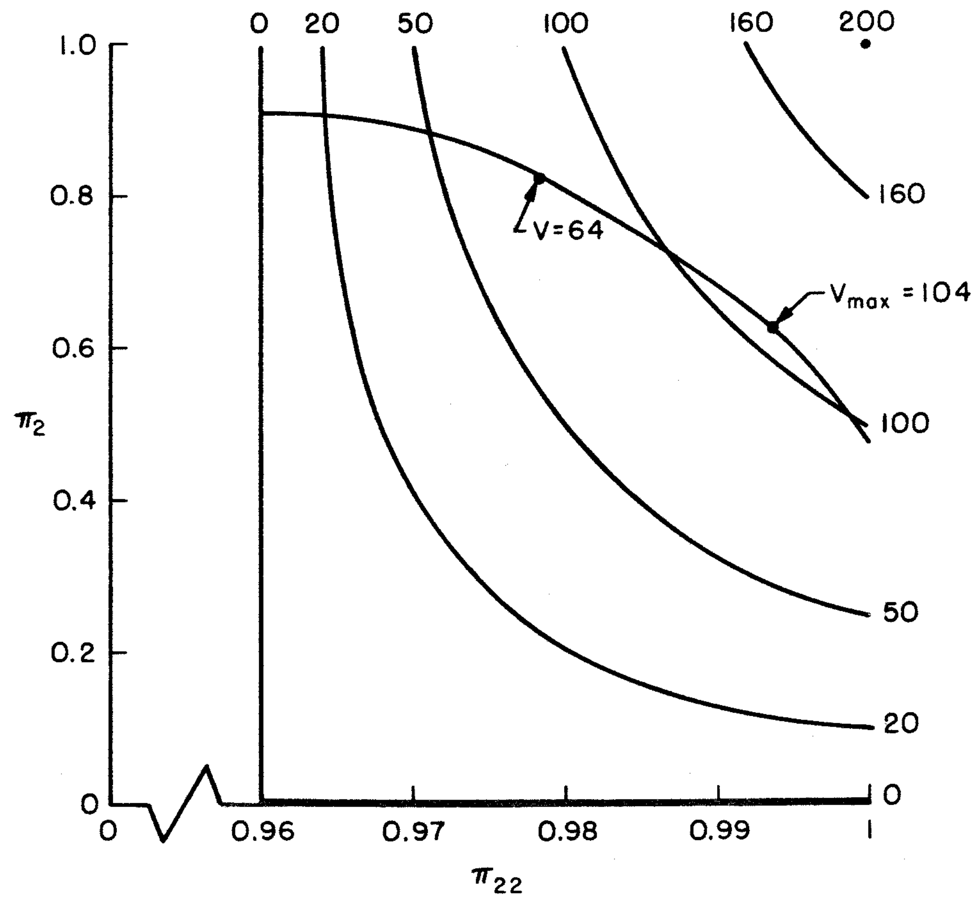
0
20
50
100
160
200
1.0
0.8
0.6
0.4
0.2
0
V=64
V_{max} =104
160
100
50
20
0
π_2
0
0.96
0.97
0.98
0.99
1
π_{22}

Fig. 11

the highest value of π_{22} achievable, consistent with a given value of π_2 and the given state of knowledge of weather phenomena summarized in the objective forecasting scheme.[1]

For the best forecast structure $\pi_2 = .63$, and $\pi_{22} = .994$. Since the climatological probability of rain is .09, π_{11} must be .232. Using the notation of the previous sections, we characterize the optimum forecast structure by the information vector

---

[1] Let us call a forecast structure characterized by accuracies (π_{11}, π_{22}) efficient if there is no other achievable forecast structure with accuracies (π_{11}', π_{22}') such that $\pi_{11}' \geq \pi_{11}$ and $\pi_{22}' \geq \pi_{22}$, with the strict inequality holding at least once. That is, a forecast structure is efficient if there is no other achievable forecast structure which is at least as accurate for both forecasts and more accurate for at least one forecast. Now consider a single forecast, say the forecast F_2 for good weather. One of the consistency properties among the probabilities we are dealing with is

$$\pi_1 \pi_{11} + \pi_2 \pi_{12} = P_1 ,$$

or

$$(1 - \pi_2) \pi_{11} + \pi_2(1 - \pi_{22}) = P_1 ,$$

where P_1 is the climatological probability of rain. Treating π_2 as a function of π_{11} and π_{22}, and the differentiating partially with respect to these two variables, we find

$$\frac{\partial \pi_2}{\partial \pi_{11}} = \frac{\pi_{22} + P_1 - 1}{(1 - \pi_{11} - \pi_{22})^2} > 0 , \text{ and}$$

$$\frac{\partial \pi_2}{\partial \pi_{22}} = \frac{\pi_{11} - P_1}{(1 - \pi_{11} - \pi_{22})^2} > 0 .$$

Hence, if it is possible to find a forecast structure which increases π_{11} with π_{22} held constant, or π_{22} with π_{11} held constant, or to increase both simultaneously, it is possible to find one which increase π_2 and π_{22} simultaneously. It follows that all forecast structures corresponding to points on the downward sloping portion of the locus in Fig. 11 are efficient, since if they are not, there would be achievable points above and to the right of that locus, which there are not. (The set of achievable points includes not only the points on the locus but also points corresponding to any other division of the Y_2 axis into intervals in which f_1 is made and intervals in which f_2 is made.)

$$I_f^* = [.37, .63; .232, .768; .006, .994] .$$

The information vector we considered in Section I was (in decimal form);

$$I_f = [.18, .82, .389, .611, .024, .976] .$$

Comparing the two information vectors, we can see that the optimum forecast structure yields the no-rain forecast less often but with higher accuracy, while the accuracy of the rain forecast is lower (less than .25).

It is interesting to note that a categorical forecast of the weather most likely to occur would have no value to this decision maker whatsoever. From Fig. 7 we see that the value of Y_2 corresponding to a probability of rain of .5 is 71. If we take this as Y_2^* (forecasting rain whenever its probability exceeds .5), we find from Fig. 9 that the resulting value of π_{21} is .05, which is greater than the critical value .04 = C/L.

As the second example of the determination of optimum forecasts based on the Thompson objective scheme, we will consider very briefly the problem of forecasting for the decision maker in the movie industry two-way call example of Section II. In this case any amount of rainfall greater than zero is unfavorable weather, so the first step is to modify Table 1 to correspond to the new definition of rain. This is done in Table 2. The relation between Y_2 and the probability of rain is shown graphically in Fig. 12. If the probability of rain is P_1, the expected returns to actions a_1, a_2, and a_3 are

$$R_1 = 0$$
$$R_2 = P_1 (- 16{,}000) + (1 - P_1) \, 9000 , \quad \text{and}$$
$$R_3 = P_1 (- 3000) + (1 - P_1) \, 5000 .$$

Solving the inequalities $R_1 > R_3$ and $R_3 > R_2$, we find that the decision

Table 2

Y_2 interval	Number of Days	Days of rain (Rain > 0)	Relative frequency
Less than 0	54	0	.00
0 to 10	55	1	.02
10 to 20	60	2	.03
20 to 30	45	3	.07
30 to 40	54	10	.19
40 to 50	36	6	.17
50 to 60	21	12	.57
60 to 70	14	9	.64
70 to 80	9	7	.78
Over 80	10	10	1.00
Total	358	60	.17

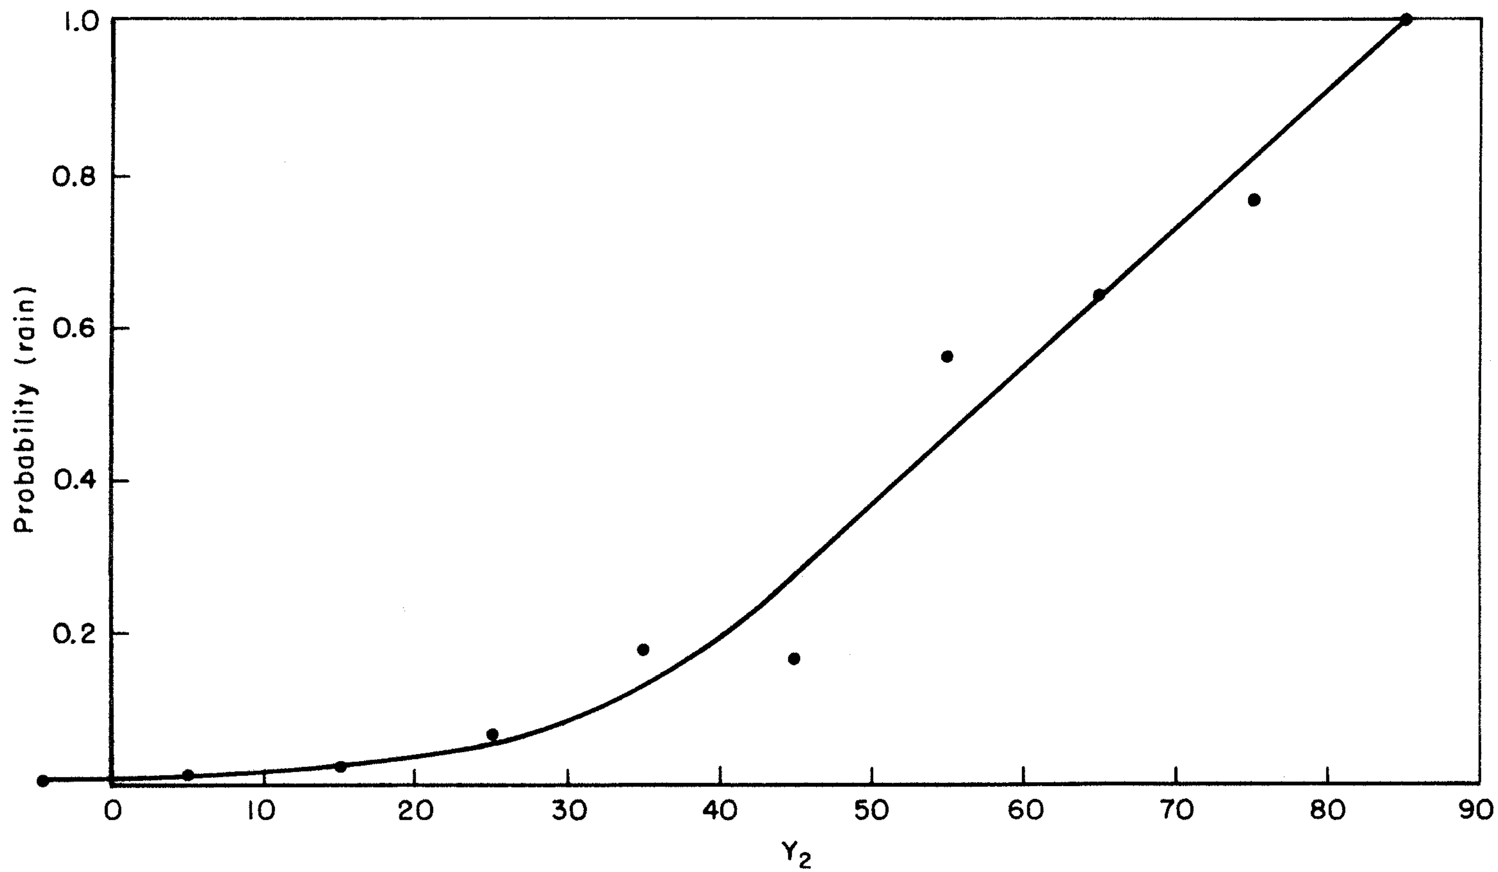

Fig. 12

maker should take action a_1 if P_1 is greater than .625, action a_2 if P_1 is less than .245, and action a_3 if P_1 is between these two values. The values of Y_2 corresponding to $P_1 = .245$ is 43, and the value corresponding to $P_1 = .625$ is 65. Hence, if the forecaster makes forecast f_1 (the high confidence forecast of rain) whenever Y_2 is greater than 65, forecast f_2 when Y_2 is between 43 and 65, and forecast f_3 when Y_2 is less than 43, he will be providing the decision maker with a set of sufficient signals. The resulting frequencies of the three forecasts and the conditional probabilities of rain associated with the forecasts are as follows:

$$\pi_1 = \frac{22}{358} = .06 \qquad \pi_{11} = \frac{19}{22} = .86$$

$$\pi_2 = \frac{57}{358} = .16 \qquad \pi_{12} = \frac{24}{52} = .42$$

$$\pi_3 = \frac{279}{358} = .78 \qquad \pi_{13} = \frac{17}{279} = .06$$

These figures are derived, of course, from the data on the frequency distribution of Y_2 and the relative frequency of rain associated with different values of Y_2 which is summarized in Table 2.[1] The expected revenue of the decision maker when employing this set of forecasts in the optimal way is

[1] In addition to the figures in Table 2, the following data on the distribution in the intervals 40 to 50 and 50 to 60 (from Thompson's Figure 5, p. 118, op. cit.) is required:

Y_2 interval	Number of days	Days of rain
40 to 43	11	1
60 to 65	11	7

$$R^*(I_f) = \pi_1(0) + \pi_2[\pi_{12}(-3000) + (1 - \pi_{12})\ 5000]$$
$$+ \pi_3(\pi_{13}(-16{,}000) + (1 - \pi_{13})\ 9000]$$
$$R^*(I_f) = .16\ (1640) + .78\ (7500) = 6112\ .$$

As shown in Section II, the expected revenue of the decision maker if he acted optimally on the basis of climatological information alone would be .17 (-16,000) - .83 (9000) = 4750. Hence, the value of the optimal forecast based on the objective scheme is 6112 - 4750 = 1362. The optimum forecast is worth more than the forecast structure considered in Section II, which had a value of 1210. In the optimum forecast structure, forecast f_3 is made more often and forecast f_1 is made less often than in the forecast structure of Section II. And the f_3 forecast is somewhat less accurate. It is also of interest to compare the value of the optimal forecast achievable using the objective scheme with the value of a forecast which simply indicates the type of weather most likely to occur. Let f_1' be the forecast of rain and f_3' be the forecast of favorable weather made in this way. The value of Y_2^* is 57, since this is the value of Y_2 corresponding to $P_1 = .5$. The resulting frequencies of the forecasts and conditional probabilities of rain are:

$$\pi_1' = .11,\ \pi_{11}' = .75\ , \quad \text{and}$$
$$\pi_3' = .89,\ \pi_{13}' = .09\ \ .$$

The conditional probabilities of rain tell us that when forecast f_1' is made, action a_1 should be taken, and when f_3' is made, action a_2 should be taken. The expected revenue is

$$R^*(I_f') = .89\ [.09\ (-16{,}000) + .91\ (9000)] = 6008\ .$$

In this case, the "most likely alternative" forecast is worth 1258 a day, which is 92 per cent of the value of the optimal forecast.

Our two examples illustrate the fact that the simplest set of sufficient signals generally involves a number of distinct signals equal to the number of actions the decision maker can take. In the next section, we examine this relationship more closely and assess its implications for the problem of labeling the forecasts in such a way as to assure that the signals given are both sufficient and intelligible.

MULTIPLE ACTIONS, MULTIPLE USERS, AND PROBABILITY FORECASTS

In this section we shall show that when forecasts are made for a user with a large number of possible actions, or when the forecasts are made for a large number of possible users, then from the point of view of the consumers of the forecasting service, it is desirable that there be a large number of different forecast signals. And this strongly suggests forecasts made in probability terms. Probability forecasts not only have the merit of being more "honest." They also have the advantage of communicating the relevant information to the decision maker in the simplest possible way. The language of probability is the most intelligible language for precise weather forecasts.

Consider one decision maker who must choose among m possible actions. We shall assume that there are only two relevant weather states -- say, rain and no rain. We shall further assume that the forecasting scheme permits rather fine discrimination of the probability of rain, and that it yields objectively correct values of the probabilities. That is, when the scheme declares the probability of rain is P_1', the historical relative frequency of rain is P_1'.

When the probability of rain is P_1, the decision maker's expected cost, using action a_i, is $C_i = P_1 C_{i_1} + (1 - P_1) C_{i_2}$. By solving the inequality $C_i < C_j$ we find the value of P_1 for which the expected cost using the ith action is less than the expected cost of using action j. By repeating this calculation for all $j \neq i$ and by finding the common part of all these intervals, we can find the set of values of P_1 for which action i minimizes expected cost.

By repeating this operation for every i from 1 to m, we can find at most m disjoint intervals covering the range of P_1 such that in each interval a particular action minimizes the expected cost. These intervals are bounded by at most $m - 1$ critical points within the range of P_1.[1] This partitioning defines, of course, an optimal forecast structure from the point of view of the decision maker in the sense that, if a particular forecast is associated with each of the intervals so defined, the decision maker's cost is the same as it would be if he had complete knowledge of the probability P_1 of rain. There is no other forecast structure involving an equal or smaller number of forecasts for which this is true.

Note that for a decision maker with m possible actions,[2] any forecast structure which is capable of making fine distinctions among probabilities of rain and which makes less than m distinctions (provides less than m

---

[1] If we assume that each action of the decision maker is optimal for some set of probabilities, then there will be exactly m disjoint regions and $m - 1$ critical points. However, the number of critical points will be less than $m - 1$ if for some action i there is another action j which is better, no matter what the weather is, or if the common part of the intervals for which it is better than j for all j, is null.

[2] And assuming that each of these actions is optimal for some set of probabilities.

different signals each associated with a different probability) is less than optimal. Conversely, if the forecasting service chooses the best mapping, it can convey all the information the decision maker can use with just m different signals.

Now assume that there are N decision makers. If individual k has m_k possible actions, and if each action is optimal for some set of probabilities, we can derive a set of at most

$$\sum_{k=1}^{N} (m_k - 1)$$

distinct critical points,[1] and these points divide the range of P_1 into at most

$$M = \sum_{k=1}^{N} (m_k - 1) + 1$$

intervals. To provide the N users of the forecasting service with information which tells each of them what he should do, we need to assign a particular forecast to each of these intervals. If the forecasting scheme is capable of making such a fine distinction among probabilities and less than M different signals are used by the forecaster, he is providing his customers with less information than they can effectively use. If it is no more confusing or costly to use M different signals than a fewer number, there is no excuse for not using M signals. For such systems as the Thompson Objective Scheme, which does permit a fairly fine partitioning of P_1, the choice of a forecast structure which denies relevant information

---

[1] If for each decision maker, each action is optimal for some probabilities, and if considering all N decision makers none of the critical points is the same, then the number of critical intervals exactly equals M .

to the decision makers requires a cost justification.

If, for some reason, the language of probability cannot be used in the forecast, then such a justification exists. Unless the critical points for the various decision makers happen to be closely grouped around relatively few values, as might be the case if they were all engaged in similar activities, one would expect that the number $M' \leq M$ of distinct forecasts required would be quite large, and that the interval of P_1 values corresponding to any particular forecast would be quite small. If the language of probability is not used, the question arises as to how a large number of forecasts can be named so as to be distinguishable and intelligible. One can go only so far in the direction of "very very slight chance of rain, very slight chance of rain, rather slight chance of rain, ... rain very likely, rain very very likely" before the whole situation becomes rather confused.

But if forecasts can be expressed in probability terms, this objection is no longer relevant. The case for probability forecasts becomes even stronger when it is noted that our preceding analysis has assumed that the forecaster is able to assign a particular forecast signal to each of the

$$\sum_{k=1}^{N} (m_k - 1) + 1$$

probability regions, and this implies that he knows what these regions are; that is, the forecaster has a very detailed knowledge of each consumer's cost matrix. If he does not know these regions, then

$$\sum_{k=1}^{N} (m_k - 1) + 1$$

different signals will not do the job. However, the use of probabilities

as the language of forecasts relieves the forecaster of the necessity of having a very detailed knowledge of the decision problems of the consumers of the forecast, and it makes the forecast very much more intelligible than it would be if an attempt were made to approach the same level of precision with forecasts quoted in categorical terms. For the decision maker to use a forecast properly, he must know its probability meaning. If the forecast is made explicitly in probability terms, there is no need to publish a guide to the interpretation of the forecasts to give the decision maker an idea of the probability of rain associated with the forecast. Each decision maker can solve his own problem and determine his own critical points; there is no need to centralize all of this information in the hands of the forecaster.

The arguments for probability forecasts would, of course, be seriously weakened if a forecast structure involving probability forecasts were considerably more expensive, or were less intelligible to the user, than a forecast structure which did not use signals in probability language. Casual observation leads us to believe that the costs of transmitting probability forecasts to a large number of decision makers are definitely negligible relative to the potential benefits from supplying them with more detailed information. A probability forecast does not require a great deal more space in the newspaper than a categorical forecast; it probably requires less. But there remains the costs of obtaining the information in probability form in the first place, and the costs to decision makers of interpreting the probability forecasts. Undoubtedly, the most important consideration here is how convenient or inconvenient it is to forecasters and decision makers to think in probability terms. It is often argued that people resist having predictions made in probability terms; they prefer a simple "it will"

or "it won't," or at least a verbal description of the degree of uncertainty that exists ("slight chance," "quite likely," etc.). Perhaps most forecasters are so used to thinking in categorical terms that they also would resist making their forecasts in probability terms.[1] We do not know to what extent this resistance represents a habit of thought which can quickly be overcome, and to what extent it represents a real psychological cost which should be taken into account in judging the merits of alternative forecast structures. We lean toward the former explanation. Also, it should be recognized that the common practice of presenting categorical forecasts and never presenting any historical data on the relationship between the forecasts and the state of the weather actually observed leaves the decison maker -- who may be able to employ probability forecasts and who may know the critical points for his various decisions very well -- saddled with the job of discovering the relationship between the categories used in the forecasts and actual weather events. This job is one which the forecaster is much better equipped to perform, and which he probably performs anyway. There is no reason why he should not make his results available to decision makers. It will not come as a revelation to the decision maker that forecasts are less than perfectly correlated with observed events, and the information will help him with his decision problem.

---

[1] In the case of an objective forecasting scheme like that considered above, there is, of course, no difficulty in presenting the forecast in probability terms, and there is a tendency for the probability estimates to be more reliable (closer to the historical relative frequency of weather states) than in the case of subjective forecasts. This is one of the advantages of the objective scheme which Thompson stressed very heavily in his article (*op. cit.*). He also shows how (in our terms) an inappropriate design of the forecast structure at the stage where the value of Y_2 is translated into a forecast can result in the loss of the economic benefit potentially realizable from the objective scheme.

In this section, we have barely touched the surface of the problem of optimal design of forecast structures to exploit a given level of knowledge of weather phenomena. We have not considered, for example, the problems arising from the multi-dimensional character of the weather, nor have we considered the case where the decision maker's payoff depends, e.g., on the amount of rain, rather than merely on its occurrence or non-occurrence. Nor have we tackled the difficult case where the payoffs of different decision makers receiving the same forecast interact, which may result in a divergence between the social and private value of the forecasts. Clearly, there is also a great deal more that could be done in the way of empirical determination of the costs of acquiring, transmitting, and interpreting weather information. In future work, we hope to consider some of these questions more thoroughly, and also to extend our analysis to cover the design of optimal forecast structures in the case where decisions are made sequentially.

Although we have emphasized throughout that we are holding the level of knowledge of weather phenomena constant while exploring the alternative forecast structures that are available, it should be pointed out that this question of forecast structure design has direct relevance to the problem of valuing improvements in the understanding of weather phenomena, and hence to the optimal use, for example, of types of observations which were hitherto unavailable. Without some knowledge of the forecast structure to be employed, there is no way of answering the question of what an improved understanding of weather phenomena might be worth. And we should not confuse the value of such improved understanding with the value of a change in the forecast structure for a given level of understanding. The most

significant comparison is between the values of weather information at two levels of understanding when each level of understanding is exploited by an optimal forecast structure.

V. THE DIRECTION OF FUTURE RESEARCH

The work presented in Sections I through IV should be considered a pilot study. We believe that it demonstrates the usefulness of decision theory as a framework for empirical research on the value of meteorological information. The work we have done falls into roughly three major parts. First, we developed a general analysis of the value of the forecasting service of a given quality to a single decision maker. We showed how the value of the forecasting service depends on the economic characteristics of the decision problems, and on the qualities of the forecast. Our analysis permitted us to value any particular improvement in weather forecasting to a particular decision maker. Second, we modified this general framework in order to provide more convenient tools for dealing with "delay" problems. Third, we examined the problem of how a weather forecaster can exploit his understanding of weather phenomena to make forecasts of the greatest possible value. We studied a number of empirical examples within our theoretical framework.

Our future plans call for an extension of an analysis of sequential problems, and empirical studies of the value of weather forecasts in situations where we suspect their significance may be considerable. We intend to look at raisin growing and air-cargo operations. But the major part of our empirical research will be on storm warnings.

Storm warnings must be studied within a sequential decision framework. The reliability of a forecast that a storm will hit increases as time goes by and the storm draws nearer, but if the decision maker waits for more unequivocal evidence of the danger, he may forgo some opportunities for

taking protective actions, or saddle himself with the higher cost of taking such measures in a hurry if the storm becomes imminent. Having more reliable forecasts -- that is, greater reliability at an earlier time -- eases the burden of this unpleasant choice between taking precautions early and perhaps incurring the costs needlessly, or deferring the precautions so long that they are unavailable or extremely expensive if actually needed. We will attempt a thorough analysis of this decision problem, and we hope to apply it in estimating the benefit of improved hurricane warnings.

The general structure of the hurricane protection problem is not hard to write down formally. Instead of considering individual actions, forecasts, and weather states as in Sections I and II, we consider sequences of actions, forecasts, and weather states occurring over some relevant period of time: for example, from the time that the existence of a particular storm is first learned of to the time when the danger from the storm is known to have ended. Formally, let $a(t)$ be the action taken at time t, $w(t)$ the weather state at time t, and $f(t)$ the forecast received at time t. In an m period problem, let vectors A, W, and F be defined as follows: $A = [a(1), a(2), \ldots, a(m)]$; $W = [w(1), w(2), \ldots, w(m)]$; $F = [f(1), f(2), \ldots, f(m)]$. The payoff or cost depends on the sequence of m actions and the sequence of m weather states: $C = C(A, W)$. Similarly, I_F is generalized into the joint distribution of sequences of forecasts and weather states. Then the problem is to choose a rule relating $a(t)$ to $f(t)$, $w(t)$, and actions and weather states occurring up to t in such a way as to minimize expected cost, given I_F. This formal characterization is

straightforward, but actual computation of optimum rules can be very difficult where the numbers of actions, weather states, and forecasts are large and the interaction of these things through time is complex.

However, note that the delay problem with penalty (treated in Section III) was of the type sketched above. We were able to take advantage of certain special structural aspects of the problem to develop reasonably simple computational techniques. It is our hope that we will be able to find similar techniques for the storm-warning problem.

The potential payoff to improvements to our hurricane-warning system, in terms of lives saved and property damage averted, may be very great. Moreover, it appears that some improvements may come as a result of cloud cover observations from satellites. The ability of a weather reconnaissance satellite to detect tropical storms has been dramatically demonstrated by TIROS. The extension of the satellite techniques to provide a careful and continuous monitoring service over the generating areas of tropical storms is possible with methods and equipment which require no new inventions or discoveries. Thus we can assume that frequent and accurate position data on tropical storms will be a direct output of meteorological satellites. This position data, and studies of the accompanying cloud system, will give the meteorologist new knowledge which may contribute to better and hence more valuable storm advisories.

Two things should be stressed. First, it is not at all clear that the impact of improved position data alone will be particularly striking in terms of economic benefits. However, this is the only contribution of satellites to hurricane warning systems that we presently feel able to treat in any systematic way. Second, the contribution of satellite

observation systems undoubtedly will involve much more than just position data. Although it is now extremely difficult to predict what improvements in storm advisories eventually will result from satellites, our economic analysis should give us some feel for where the big payoffs are, and a rough estimate of their value.

Weather satellites will have value if they yield more or better information than existing systems for a given cost, or if they yield roughly comparable information at a smaller cost. Although our proposed analysis certainly will not yield an exact figure for the value of satellite observations in hurricance-warning systems, we hope that the results will provide some basis for future policy.

APPENDIX 1 TO SECTION I

THE COMPUTATIONAL FRAMEWORK

Let there be n possible decisions, $a_1 \ldots, a_i \ldots, a_n$, and m relevant weather states, $w_1 \ldots, w_j \ldots, w_m$. Let c_{ij} be the cost incurred by the decision maker if decision a_i is taken and weather w_j occurs. Thus the following cost matrix is defined:

$$\begin{array}{c|ccccc} & w_1 & \cdots & w_j & \cdots & w_m \\ \hline a_1 & c_{11} & \cdots & c_{1j} & \cdots & c_{1m} \\ \vdots & \vdots & & \vdots & & \vdots \\ a_i & c_{i1} & & c_{ij} & & c_{im} \\ \vdots & & & & & \\ a_n & c_{n1} & & c_{nj} & & c_{nm} \end{array} \qquad (1a)$$

$$c = c(a, w) \qquad (1b)$$

Let the climatological probabilities of the various weather states be: $I_o = (P_1, P_2 \ldots, P_m)$ such that $\sum_{j=1}^{m} P_j = 1$. If the decision maker must choose among his alternative actions on the basis of climatological information alone, we shall assume that he should choose the action that minimizes expected cost. He should choose action a^* such that:

$$C^*(I_o) = \text{Min}\ (C_1, C_2 \ldots, C_i \ldots, C_n), \qquad (2)$$

where $C_i = \sum_{j=1}^{m} C_{ij} P_j$ is the expected cost if action a_i is taken.

Assume that there is a forecasting structure with q different possible forecasts: $f_1, f_2 \ldots, f_k \ldots, f_q$. Each forecast defines a conditional probability distribution of weather:

$$P(W|f_k) = (\pi_{1k}, \pi_{2k} \ldots, \pi_{jk} \ldots, \pi_{mk}) ,$$

where π_{jk} is the conditional probability that weather state w_j will occur, given that forecast f_k has been received. Clearly $\sum_{j=1}^{m} \pi_{jk} = 1$ for $k = 1 \ldots, q$. The relative frequency of each forecast f_k is π_k, such that $\sum_{k=1}^{q} \pi_k = 1$. We know that the π_k's and the $P(W|f_k)$'s must be related to the climatological probabilities as follows:

$$\pi_{11}\pi_1 + \pi_{12}\pi_2 + \ldots + \pi_{1q}\pi_q = P_1 , \quad \text{and} \tag{3a}$$

$$\pi_{21}\pi_1 + \pi_{22}\pi_2 + \ldots + \pi_{2q}\pi_q = P_2$$

$$\cdot \qquad \cdot \qquad \qquad \cdot \qquad \cdot$$

$$\pi_{n1}\pi_1 + \pi_{n2}\pi_2 + \ldots + \pi_{nq}\pi_q = P_q .$$

$$\sum_{k=1}^{q} \pi_{jk}\pi_k = P_j . \tag{3b}$$

Let us define:

$$I_f = [\pi_1, \pi_2 \ldots, \pi_q; P(W|f_1), P(W|f_2) \ldots P(W|f_k) \ldots P(W|f_q)] . \tag{4}$$

I_f is derivable from a $q \times m$ contingency table. If the decision maker chooses the action to minimize expected cost, let $C^*(f_k)$ be the expected cost when the forecast is f_k.

$$C^*(f_k) = \text{Min}\,[C_{1k}, C_{2k} \ldots, C_{ik} \ldots, C_{nk}] \quad \text{for} \quad k = 1 \ldots q , \tag{5}$$

where $C_{ik} = \sum_{j=1}^{m} C_{ij}\,\pi_{jk}$ is the expected cost if action a_i is taken,

given that forecast f_k has been received. The $C^*(f_k)$'s define the "best" action function:

$$a = \hat{\alpha}(f) \quad . \tag{6}$$

If the decision maker follows this best decision rule, his expected cost will be:

$$C^*(I_f) = \pi_1 C^*(f_1) + \pi_2 C^*(f_2) + \ldots + \pi_k C^*(f_k) + \ldots + \pi_q C^*(f_q) \quad . \tag{7}$$

Thus $C^*(I_f)$ is the decision maker's expected cost if he has information I_f and makes his decisions so as to minimize expected cost.

It can be shown that $C^*(I_f) \leq C^*(I_o)$. One demonstration of this fact is that I_o can be deduced from I_f ; thus, at worst, the decision maker can use I_o instead of I_f .

We define the value of I_f as:

$$V(I_f) = C^*(I_o) - C^*(I_f) \; , \tag{8}$$

and we only need matrix and vector (4) to compute (8).

To set an upper bound on the value of any improvements in meteorological information we can compare $C^*(I_f)$ with $C^*(I_{oo})$, the average cost of decision maker would incur if he could predict weather perfectly. Clearly for I_{oo} the number of forecasts, q , should equal the number of relevant weather states, m , and for each forecast, f_k , $\pi_{kk} = 1$ and π_{ij}, $j \neq k = 0$.

Thus:

$$I_{oo} = (P_1, P_2 \ldots P_m; 1, 0 \ldots 0; 0, 1, 0 \ldots 0; \ldots; 0 \ldots, 1) \tag{9}$$

The calculation of $C^*(I_{oo})$ is straightforward.

APPENDIX 2 TO SECTION I

SOME ISSUES RELATING TO CLIMATOLOGY AND PERSISTENCE

CLIMATOLOGY

In Section I we defined climatological information as long-run relative frequencies of different relevant weather states at different times of the year and in different locations. In terms of the general model (sketched out under pp. 21-23 of this section) in which we modeled the forecaster's problem as that of choosing an optimum η for the equation $f = \eta(w, z, t, L)$, climatological predictions can be expressed as $f = \eta(t, L)$. The η defines a probability distribution of weather states as a function of time and place. The η of course can only be as good as the underlying climatological data.

Even our reasonably precise definition of climatology is not unambiguous. For example, is the "best" climatological probability of rain on July 4 in Los Angeles the percentage of past July 4ths on which it has rained? Is it the proportion of days during the first week of July on which it has rained? Or is it a longer run moving average? Clearly the "true" climatological probabilities change as t changes, but how rapidly and how erratically? These problems would not be serious if we had very many years of observation, but they may be important when we do not.

We assume that good data are available in sufficient quantity so that the calculated climatological probabilities are the "true" ones. We shall not discuss this assumption in any detail here, except to raise the question of how much "true" climatological data are worth. Remember, our $V(I_f)$ is calculated as value over "true" climatology.

Perhaps we can place an upper bound on the value of climatological information by assuming that, in the absence of such information, the decision maker operating in complete ignorance would choose the same action day after day, and further choose the worst possible action. We could then calculate how well the decision maker would do if he used this "worst possible" strategy, in comparison with how he would do if he had good climatological data. The difference would be an upper bound to the value of climatological information. And of course we need climatological data to do these calculations.

In our trucking example we found that the best day-after-day action for the decision maker was a_1 . His average cost would be \$200 a day using a_1 , while action a_2 has an average cost of \$450 a day. Thus 450 - 200 = 250 seems definitely to be an upper bound on the value of climatological information.

Perhaps a more reasonable number might be obtained if it is assumed that in the absence of any knowledge of probabilities, the decision maker would "minimax." That is, he would pick the action which, at its worst, would yield a lower cost than any other action at its worst. Thus in the trucking example if action a_1 is taken, cost cannot exceed \$200 no matter what the weather turns out to be, while if action a_2 is taken cost will be \$5,000 if weather w_1 occurs. Thus a_1 is the minimax strategy. But notice that for this particular example if $C^*(I_0)$ is compared with $C^*(\text{Minimax})$, the value of climatological information is zero. Different numbers in the cost matrix would, of course, yield a different result.

Many other "complete ignorance" strategies could be assumed and the value of climatological information calculated accordingly. The "worst

possible single action" seems a reasonable upper bound assumption, but climatological information is almost certainly worth less than this upper bound. And clearly the "complete ignorance" assumption as an alternative to "climatological information" may not be particularly relevant to any interesting questions. One would guess that anyone who operates in an area for a significant period of time would develop a <u>feel</u> for climatological probabilities that might not be too far wrong.

PERSISTENCE AND HOME-MADE FORECASTS

In calculating the value of forecasts we have been assuming that, in the absence of forecasts, the best predictions a decision maker can make are climatological predictions. But the individual decision maker often has access to meteorological information of wide variety and from several different sources. In particular, in addition to using formal or informal climatological data, the weather forecast in the daily paper or over the radio, or a specially prepared forecast, the decision maker can use his own observations of many weather variables. It is important for us to note that the decision maker, by using his own observations, can often do better than using climatology alone. Sometimes he can even do better than using forecasts alone. Ideally our analysis should take this into account.

Since weather states tend to persist, knowledge of the existing state of the weather provides a good deal of information about what the weather will be in the near future. Indeed, one of the major building blocks of most weather forecasts is persistence. It seems worthwhile to see how well the decision maker of our example would do if he based his decisions on his own observations. But first we must define what we mean by persistence.

In the case where there are only two relevant weather states, w_1 and

w_2 , persistence is easily defined. Let P_{22} and P_{11} be the probabilities that a day of good weather will be followed by another day of good weather, and a day of bad weather by another such day. As before, let P_2 and P_1 be the climatological probabilities of good weather and bad weather. If $P_{22} > P_2$ and $P_{11} > P_1$, then we say that weather tends to persist from day to day. If $P_{21} = 1 - P_{22}$ is the probability that a day of bad weather and if P_{12} is similarly defined, the reader may then verify that if there is persistence, as defined above, then $P_{21} < P_1$ and $P_{12} < P_2$.[1]

We have defined persistence in terms of day-to-day changes, although clearly the shorter the period between the moments of time we are considering the greater will be the tendency of weather states to persist. If we consider the weather from hour to hour, P_{11} and P_{22} are both close to one (and P_{21} and P_{12} are both close to zero). If we consider persistence over a period separated by a week, P_{11} is approximately equal to P_1 and P_{22} is close to P_2 . (P_{21} also is close to P_1 and P_{12} is close to P_2 .) For our analysis we are interested in persistence over the time elapsing between when the decision must be made and when the weather state counts. We shall not extend our definition of persistence to situations involving more than two weather states at this time.

In the trucking problem of Section I, how well would the decision maker do if he based his actions on his own observations? In particular, how well would the dispatcher do if he assumed that weather states persist

---

[1] The following equations clearly must hold:

$$P_1P_{11} + P_2P_{21} = P_1$$

$$P_1P_{12} + P_2P_{22} = P_2$$

and thus followed the decision rule: tarp the trucks this evening if rain today exceeds .15 inches, but not otherwise. It turns out that this decision rule is inferior to basing decisions on climatology. The probability that tomorrow will have more than .15 inches of rain, given that today does not, is about .06. This figure is larger than the cost-loss ratio of $\frac{200}{5000}$. Thus a decision rule for which the trucks will not be tarped when rain is less than .15 inches will be more costly, on the average, than a decision rule which causes the trucks to be tarped every night. Thus a pure "persistence" forecast is worse than climatology, in this example. For different numbers in the cost matrix and different values of the relevant meteorological variables, the result might be different.

But though a pure persistence forecast is inferior to climatology in our trucking example, a slightly different home forecasting scheme is better than climatology. Consider the decision rule: tarp the trucks this evening if there has been any rain at all today. The probability that rain tomorrow will exceed .15 inches, given that it has not rained at all today, is .03. This is smaller than the cost-loss ratio. Thus on the 84 days out of 100 that it does not rain at all, the dispatcher, if he has no other information, should take no protective action. However the probability that rain tomorrow will exceed .15 inches, given that some rain has fallen today, is about .37. Thus on the 16 rainy days out of 100, protective action should be taken.

Denoting information from home-made forecasts as I_H:

$$C^*(I_H) = \frac{16}{100}\,\mathrm{Min}[200, \frac{37}{100} \times 5000] + \frac{84}{100}\,[200, \frac{3}{100} \times 5000] = 158 \ .$$

Comparing costs using professional forecasts with costs using home forecasts:

$$V(I_f) = C^*(I_H) - C^*(I_p) = 158 - 136 = 22 \ .$$

In our trucking example the possibility of making simple home forecasts on the basis of easy observations reduces the value of the forecasting service to 22 a day. In other examples it might turn out that, relative to climatology, any home forecasting scheme is worthless.

Notice that when we consider the relevant partitions of the possible weather states from the point of view of making a forecast, or providing a signal for a decision rule, this partition may not coincide with the partition relevant to the cost matrix. If our dispatcher makes his own observations and bases his decisions on them, he must distinguish no rain from rain of less than .15 inches, even though this distinction does not enter the cost matrix. Similarly, the extent of cloudiness, the wind direction, etc. may provide useful inputs to a forecast, even though they are not directly relevant to cost. This is why we have written the forecast function $f = \eta(w, z, t, L)$. The z represents a different partition of weather states from that which is relevant for the cost matrix $C = C(a, w)$.[1]

---

[1] Theoretically we could make a sufficiently fine partition to include both the relevant cost partition and the relevant forecast partition. But in general such a framework would be more complex and no more useful.

APPENDIX TO SECTION IV

The forecast structures studied in this section are based on the Thompson objective system for forecasting rainfall in the Los Angeles area.[1] Thompson's discussion is drawn on heavily in what follows. The period covered by the forecast is 10:30 A.M. of one day to 4:30 P.M. the next. All of the observations upon which the forecast is based are available by 4:30 A.M., and the amount of computation involved in reaching the forecast does not significantly delay getting the forecast out. Observations on six weather variables are used in making the forecast:

1. The altitude at which a barometric pressure of 700 millibars is observed at Oakland
2. The sea level pressure at San Francisco minus that at Los Angeles
3. The sea level pressure at San Francisco
4. The sea level pressure at Los Angeles minus that at Phoenix, Arizona
5. The wind direction at Sandberg, California
6. The temperature at the altitude at which pressure of 700 millibars is observed at Santa Maria, California

The reader is referred to Thompson's article for a discussion of the theoretical and empirical reasoning which went into this particular choice of variables. It might be mentioned, however, that several other variables have been tried and make no significant addition to the quality of the

---

[1] J. C. Thompson, "A Numerical Method for Forecasting Rainfall in the Los Angeles Area," Monthly Weather Review, July 1950, pp. 113-124.

forecasts. Furthermore, a comparison of the forecasts generated by the objective scheme with those actually made, by subjective methods, at Los Angeles showed the two forecasting methods to be of comparable accuracy. A relatively small number of variables may, therefore, adequately summarize the very large amount of information which is available, at least potentially, to the forecaster.[1] It is also interesting to note, in passing, the extent to which these basic observations would be "uninteresting and perhaps unintelligible" to the average decision maker.

The relationship between these six variables and rainfall in the Los Angeles Basin was determined in a manner which can be roughly described as follows: Observations on the six independent variables and on the amount of rainfall during the hours of the day to which the forecast refers for the winters (October-March, inclusive) of 1946-47 and 1947-48 were the basic data. The independent variables were grouped into three pairs, 1 - 2, 3 - 4, and 5 - 6. For each pair of variables, a scatter diagram was made of the observed values, with the amount of rainfall which occurred after that observation indicated beside each plotted point. Then contours of equal probability or amount of rainfall were developed on those scatter diagrams, and three new variables were defined by assigning scale values from 0 to 100 to eleven spaced contours covering the region in which the observations were found. Then the first two of these new variables were combined by the

---

[1] Of course, these six variables were selected because they were effective in summarizing a large part of the total information on conditions affecting rainfall in the Los Angeles area available to the forecaster. But the fact that it was possible to find a way to combine six variables into a forecast of quality comparable to those reached by subjective means indicates that the complexity of the forecasting problem, and the usefulness of taking a great many subtle indications into account, can be overemphasized.

same procedure into another derived variable, and this derived variable was again combined with the third to yield a single final variable, called Y_2 in the Thompson article.

Each value of Y_2 is uniquely associated with a given probability of rain, and summarizes all the information about rainfall which this particular forecasting scheme generates. These probabilities, together with the relative frequencies with which the scheme generates different Y_2 values, provide a complete description of the forecasting scheme. With these data it is possible to calculate and value the alternative forecast structures, as we have done in this section.