551.509.324.3
ISSN 0110-5965

NEW ZEALAND METEOROLOGICAL SERVICE

TECHNICAL NOTE NO.245

A PROBIT MODEL FOR FROST-FORECASTING IN CENTRAL OTAGO, NEW ZEALAND

S.W. Goulter
and
D. Vere-Jones

Issued for limited distribution by:

The Director
New Zealand Meteorological Service
P.O. Box 722
WELLINGTON

18.11.81

A PROBIT MODEL FOR FROST FORECASTING IN CENTRAL OTAGO, NEW ZEALAND

S.W. Goulter[1]
and
D. Vere-Jones[2]

Abstract

The background to frost forecasting is described, and a simple model developed which provides a regression estimate for minimum overnight temperature and converts it into a forecast for the probability of frost occurrence overnight. The performance of the model is discussed, and its implementation within an on-going forecasting operation is described. Attention is given to the question of choosing a probability threshold to maximize the performance of a skill index (Hanssen-Kuiper index).

1. Introduction

This paper outlines the development and application of a simple numerical procedure for estimating the probability of an overnight air frost from information available in the afternoon. The method is developed for a particular fruit-growing area in Central Otago, New Zealand, but could be applied to other regions with suitable modifications. It is based on a probit model which incorporates a classification according to general synoptic pattern followed by a multiple regression to predict the minimum overnight temperature.

Frost prediction is an interesting problem for the meteorologist. Because frost occurrence is dependent not only on the general atmospheric flow but also on what can be highly localized physical processes, the problem involves elements which are common to a wide range of meteorological forecasting problems. At present, probability forecasts are not generally used by the New Zealand Meteorological Service. The development of this model is part of the more general exercise of attempting to evaluate their possible effectiveness in the New Zealand context. Forecasting air frosts was chosen partly because of its intrinsic interest and importance to fruit growers, and partly because it was recognized as a problem not handled very adequately by existing procedures. The aim was to produce a simple quantitative estimate of the probability of an overnight air frost which, in the first instance, could be used by the forecaster at his discretion, and made the basis of comparisons with other forecasting techniques.

---

[1] Meteorological Office, Wellington

[2] Victoria University of Wellington

Frost forecasting overseas has tended to concentrate on the prediction of overnight minimum temperatures. The proposed method makes use of such a prediction, but supplements it with a probability statement which conveys some indication of the reliability of the frost prediction. The alternative to conveying this information in terms of a probability statement would be to convey it in terms of a standard error of prediction. As the standard errors are quite large, it was felt that the probability approach was both closer to existing forecasts and more easily understandable.

2. Background: Current Forecasting Procedure

Orcharding is a feature of life in the Central Otago region of New Zealand and has been for many years. Every year a large quantity of fruit ripens and matures, and every spring this fruit is vulnerable to frost damage. Frost fighting operations are consequently important to the grower.

The low temperatures in the region generally develop in the following way: following the passage of a low pressure trough over southern New Zealand, there is an incursion of cold air, usually with a flow from the southerly quarter. Initially skies may be cloudy and it may be windy. With the advance of the next high pressure ridge, subsidence occurs leading to clearing skies and a decrease in wind. Under these conditions further radiational cooling will occur at night. Occasionally the mere advection of freezing southerly air is sufficient to cause great damage, even during daytime, but these conditions are relatively rare.

To provide some measure of warning for their frost fighting operations overnight, growers in the locality have a seasonal request to the New Zealand Meteorological Service to provide forecasts of overnight frost there during spring. Traditionally, frost forecasting has been regarded as a difficult enterprise. Dr E. Kidson (1927) states: -

> "It is not possible for the meteorologist in Wellington to forecast frost in a particular small region with great accuracy, since he has not sufficient local information, and in any case his forecasts would become impossibly voluminous if he was to attempt to give details for each locality. But he is in a position to know what the general conditions are, and his forecast should enable the local observer to know when to be on the alert."

In line with this philosophy, present forecasts are made in general verbal terms, and must mention the words either 'frost' or 'no frost'. They are primarily forecasts of ground frost, but are to mention air frost if expected. An examination of these forecasts was made for the springs of 1975 and 1976, and compared with the actual occurrence of ground frosts at Earnscleugh. The results are summarized in Table 1.

Table 1: Accuracy of Ground Frost Forecasts as at Earnscleugh Climate Station.

1975	Forecast	
	Frost	No Frost
Frost	32	2
No Frost	12	30

1976	Forecast	
	Frost	No Frost
Frost	26	1
No Frost	14	41

Earnscleugh is the location of the Government Research Orchard, and is the frostiest orcharding location for which a climatological record is available. It will be used in the sequel as a suitable location for developing and testing the forecast procedures. The Table shows that almost all ground frosts were forecast, but also that many false forecasts were made. Some over-forecasting may be desirable on the grounds that frosts may develop at other fruit-growing locations in the area even when they do not occur at Earnscleugh. For the sequel, however, we shall take as our goal the accurate forecasting of frost events for Earnscleugh itself.

In contrast to present procedures, we shall concentrate on the forecasting of air frosts. An air frost is said to occur when the temperature, as measured in a meteorological screen 1.3 metres above the ground, falls below 0^{o} Celsius. In general, air frosts are much less well-handled at present than ground frosts. In 1975, for example, there were 16 air frosts between mid-September and the end of November, yet only 5 of 76 forecasts specifically mentioned air frosts. Since air frosts are likely to be of greater relevance than ground frosts for orchardists in the region, and since, being rarer, they are more difficult to forecast, the development of a quantitative procedure for air frost prediction was considered a worthwhile supplement to existing procedures.

3. Development of the Model

3.1 Regression Equations and Synoptic Classification

The variable taken as the ultimate subject of the predictions was the frost event at Earnscleugh, using information available to the forecaster by the middle of the previous afternoon. At present, the nearest synoptic data available is from Alexandra, about six kilometres from Earnscleugh. Mid-afternoon synoptic reports from Alexandra were also available for four years, 1968-1971; the reports ceased in 1971 but have recently been restarted with the aid of local growers. Records of air minimum temperarture are available from both Earnscleugh and Alexandra, as well as other climatological stations in the locality.

In view of the limitations of the data, it was decided to develop the equations on the basis of the 1968-1971 data, using current data to provide an independent check on performance. Because synoptic information was not available from Earnscleugh itself, a two-stage procedure was adopted: first the overnight temperature fall was predicted for Alexandra; then the expected Alexandra minimum temperature was used as the input variable for the probit analysis of frost occurrences at Earnscleugh. This hybrid procedure appeared to be more robust than trying to predict the overnight temperature fall at Earnscleugh from the Alexandra data, or trying to use the synoptic data directly as input to the probit estimation procedure. The analysis throughout was restricted to the spring period, defined here as the period September 1st - November 30th.

Many variables which, on physical grounds, were thought likely to influence the temperature fall were considered in developing the regression equations. These included:

(a) dry bulb temperature at Alexandra at 3 p.m. (degrees Celsius);

(b) dew point temperature at Alexandra at 3 p.m. (degrees Celsius);

(c) cloud cover total at Alexandra at 3 p.m. (expressed as the number of eights of the sky covered);

(d) wind as in the 3 p.m. Alexandra synoptic observation (speed, in knots; direction, positive from south or 270° true);

(e) time from beginning of spring (counting Sept 1 = 1, etc a whole number of days);

(f) an index of soil moisture (based on rain reports from Alexandra);

(g) an index of snow cover in surrounding country above Alexandra.

The overnight temperature fall examined was 3 p.m. Alexandra dry bulb air temperature minus night minimum air temperature at Alexandra.

Many approaches to this type of problem overseas have begun by taking a full data set, and then rejecting all nights where the daytime observations were considered not representative of those overnight (for example, see Veitch (1959)). The difficulty with such an approach is that, in effect, the forecaster has to make a preliminary forecast on whether or not to use the regression. The present analysis was different, in that an attempt was made to cover all situations by classifying them according to a synoptic index. It is desirable that such an index should not only characterize in a general way the synoptic environment in which a particular regression equation should be used, but should also be objective and capable of consistent use by a group of forecasters with differing personal attitudes and judgements. It is an added advantage if the method involves the forecaster actively, thereby building up his experience with the technique and ability to interpret the final numerical results.

It is clear that such requirements impose rather stringent constraints on the classifications which can be used. After many trials, an index with just two values was adopted, defined AC or C according as the curvature at the grid point, 40°S, 160°E in the Tasman Sea, possessed anticyclonic (AC) or cyclonic (C) curvature at midday on the day of the forecast. This index was chosen to take account, in a simple manner, of the large scale features of the atmospheric flow. Being well into the Tasman, the grid point avoids the complicating effect of the land on the isobaric pattern. In most flow patterns it is upstream from Central Otago. Thus in mobile patterns (usually migratory eastward) the curvature at this point might be expected to be representative of the curvature over Central Otago overnight. It is usually an adequate index also in large, slow-moving anticyclonic situations; however as the belt of high pressure moves east the index occasionally becomes unrepresentative. The classification is quite objective; checks indicate that a given sequence of classifications can be repeated with a high degree of consistency. Some examples of this classification are shown in Figure 1. Almost all the frost situations were found to occur within the AC category.

SYNOPTIC CLASSIFICATION

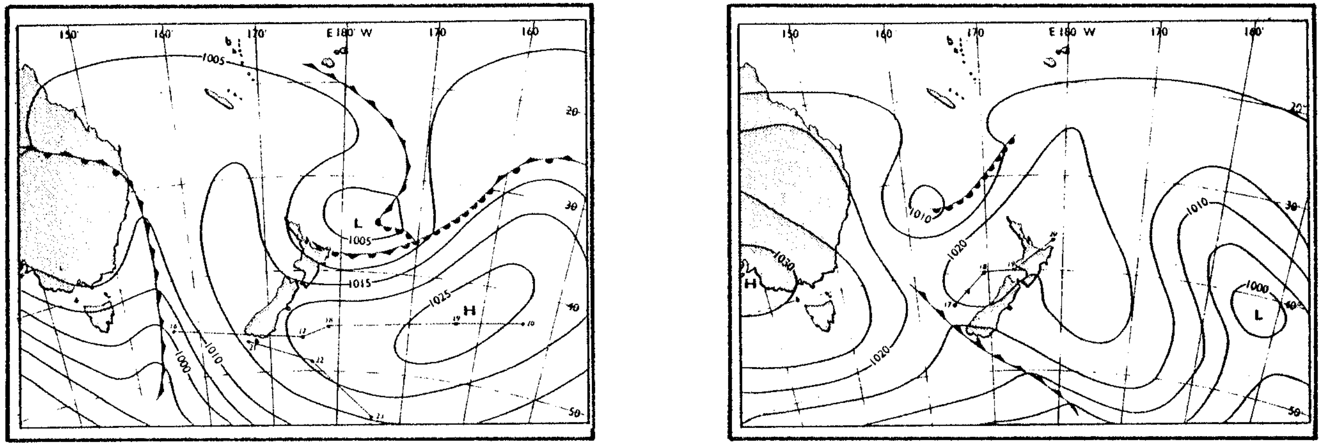

1- Anticyclonic

2- Anticyclonic

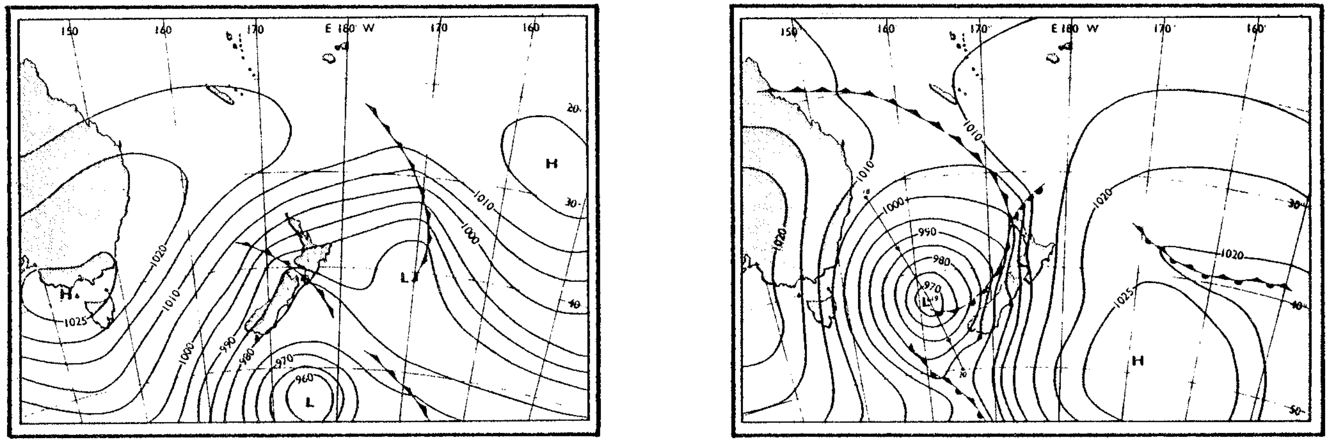

3- Cyclonic

4- Cyclonic

n‾‾‾ track of pressure system from position at day n

From: N.Z. MET. S. Misc. Pub. 115 (4)

Figure 1. Some examples of the Synoptic Classification.

This simple classification technique was adopted as a first attempt to come to grips with the full forecasting problem which faces the meteorologist. Since frosts tend to occur in well-known circumstances - i.e. in anti-cyclonic conditions where dynamical influences are relatively small - there is some hope for such a simple stratification technique, although the most difficult and interesting cases are likely to be just those where some dynamic component to the weather, such as an approaching cold front, has to be taken into account.

Regression equations were fitted within each category, using the data set (complete apart from weekend data, which were not available for the springs of the years 1968-71). Explanatory variables were restricted to information available to the forecaster at 3 p.m. The main effects were 3 p.m. temperature, cloud cover, and day number (time), the other variables contributing relatively little to the reduction of the residual standard error. The coefficients and their standard deviations, for the best-fitting equations are shown in Table 2A.

Table 2A: Coefficients of Regression Equations with Associated Standard Deviations

General Equation
Temperature drop = $\alpha + \beta \times$ temperature $+ \gamma \times$ cloud
$+ \delta \times$ time + error

Category AC	α	β	γ	δ	Standard Error	R^2
Coefficient	4.640	0.571	-0.209	-0.042	2.48	0.54
S.D.	1.043	0.054	0.100	0.009		

Category C	α	β	γ	δ	Standard Error	R^2
Coefficient	7.403	0.389	-0.483	-0.018	2.18	0.33
S.D.	1.402	0.070	0.125	0.010		

The signs of the coefficients are physically reasonable, as the temperature fall overnight should increase with higher daytime temperatures, but decrease with increasing cloud cover and time into spring.

Although the time term (number of days from Sept. 1) makes a valuable contribution in the AC category and should be retained if the main object is prediction of temperature, it was ultimately dropped in the interest of simplicity. The introduction of non-linear terms (squares or products of the explanatory variables) also effected no significant improvement to the fit. The equation carried forward to the probit analysis attempted to combine maximum simplicity with near optimal predictive power, and is shown in Table 2B.

Table 2B: Coefficients of Regression Equations with Associated Standard Deviations

Simplified equation (omitting time term)

Category AC	α	β	γ	Standard Error	R^2
Coefficient	4.851	0.454	-0.259	2.65	0.47
S.D	1.115	0.050	0.107		

Category C	α	β	γ	Standard Error	R^2
Coefficient	7.779	0.319	-0.486	2.20	0.31
S.D.	1.401	0.058	0.127		

A plot of predicted versus actual minimum overnight temperature for Alexandra is shown in Figure 2. There is a very large amount of scatter, which has been condensed by grouping together and averaging the actual temperatures corresponding to a given range of forecast temperatures. The large scatter is typical of the results obtained in such regression work on meteorological data, and reflects the wide variety of types and combinations of synoptic features which can affect a single variable such as the overnight minimum temperature.

3.2 Probit Analysis

The next stage in the development of the model was to fix an algorithm for converting the minimum temperature predictions for Alexandra into frost predictions for Earnscleugh. For this purpose a probit curve was fitted

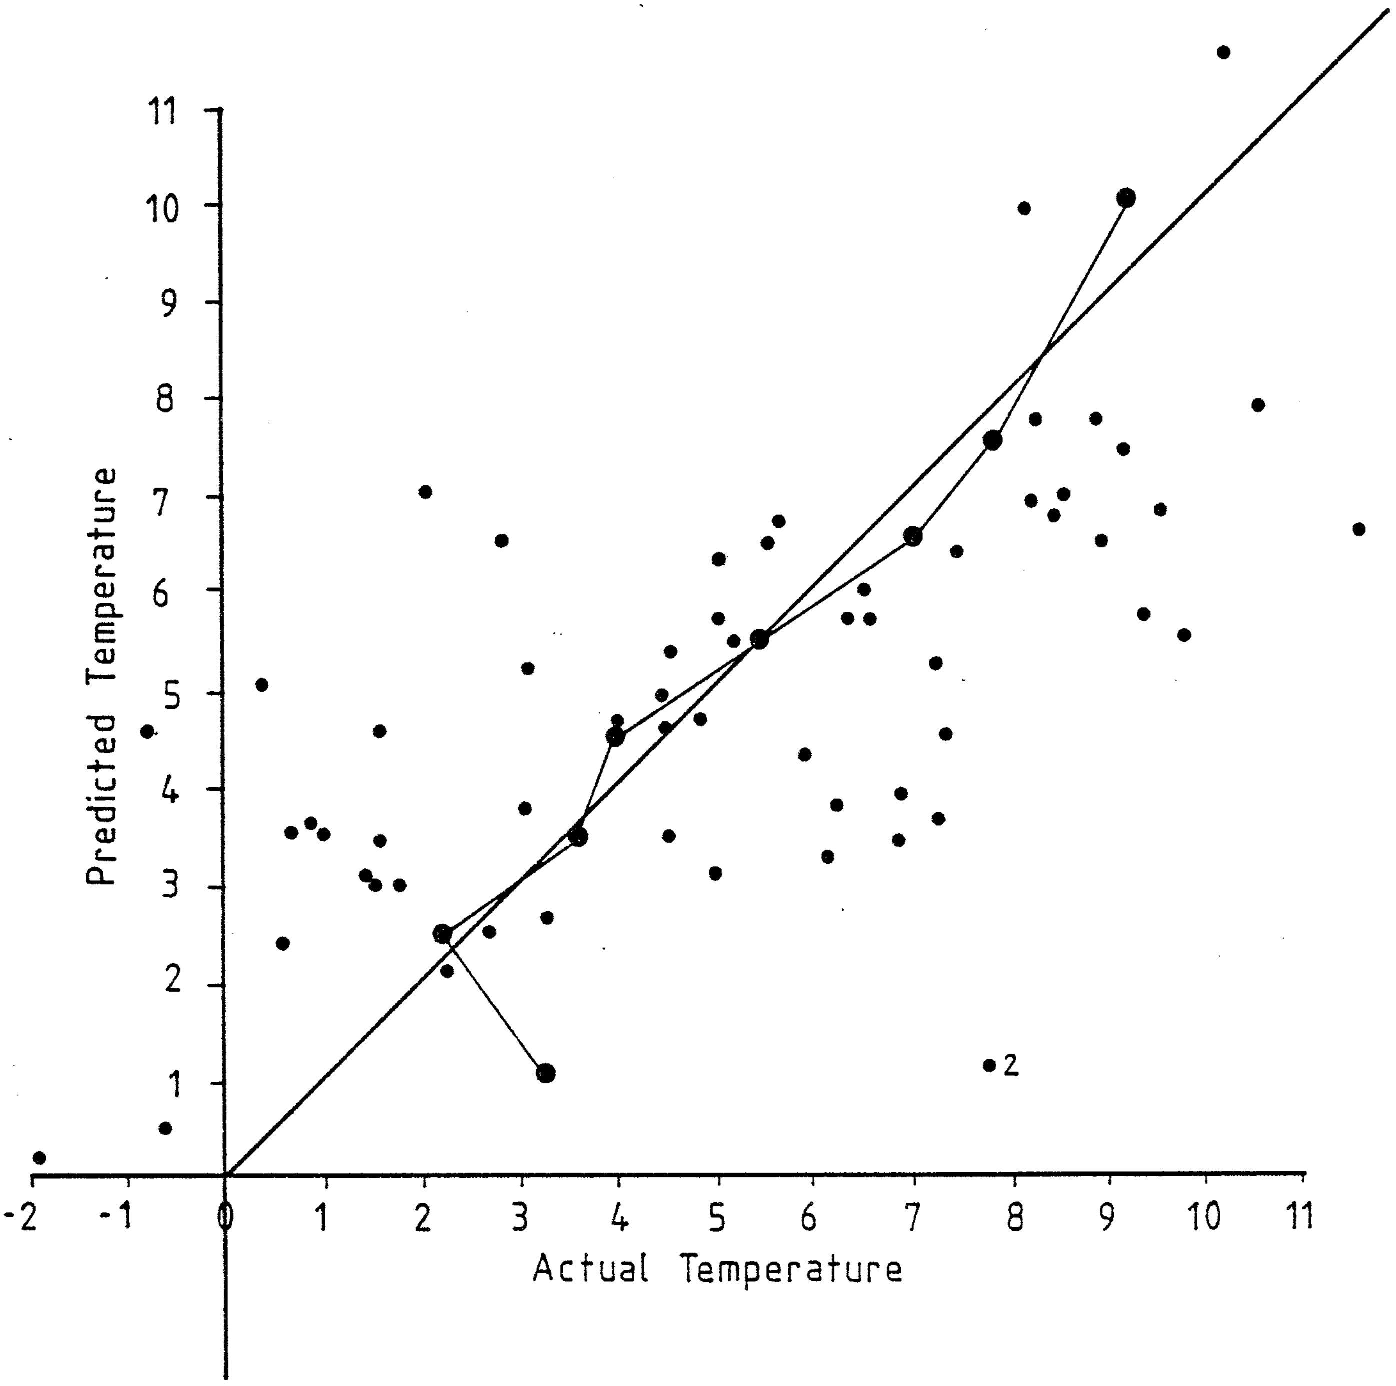

Figure 2. Predicted Versus Actual Minimum Air Temperature at Alexandra for 1968.

. values for individual nights

● averages over nights with predicted temperatures in given 1°C ranges (except for the first and last values, which are over 2° and 4° ranges respectively).

\- line y = x

to the observed occurrences of air frosts at Earnscleugh making use of the predicted overnight minimum temperature at Alexandra as an explanatory variable analogous to the "dosage" in the traditional biometric applications of the technique (see, for example, Finney (1971)). In fact both probit and logit models were tried, but the probit model gave slightly better results, and in addition is suggested by the following reasoning:

Suppose that the minimum overnight temperature at Earnscleugh (T_E) is linearly related to the minimum overnight temperature at Alexandra (T_A), say

$$T_E = k + \lambda T_A + \text{error}$$

where the errors are normally distributed with mean 0 and variance σ^2 and we might expect $\lambda \simeq 1$; then we have

$$\begin{aligned}\text{Prob}(T_E < 0 \mid T_A) &= \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{0} \exp\{-(x-k-\lambda T_A)^2/2\sigma^2\}dx \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{-(k+\lambda T_A)/\sigma} \exp(-\tfrac{1}{2}u^2)\,du\end{aligned}$$

If the critical overnight minimum temperature at Alexandra is replaced by its regression estimate $\hat{T}_A$, a further error term is introduced without, however, altering the basic form of the equation. Thus it seems reasonable to anticipate a relation of the type: -

$$\text{Prob (frost at Earnscleugh} \mid \text{forecast minimum temperature at Alexandra} = \hat{T}_A) = P\{\alpha - \beta\hat{T}_A\}$$

where P(x) is the standard normal distribution function, β is a measure of the precision of the forecast (low β corresponding to low precision and $\frac{\alpha}{\beta}$ is a measure of the temperature shift between Earnscleugh actual temperatures and Alexandra forecast temperatures).

One advantage of this two-stage approach is that only one predictor variable (the forecast temperature) has to be fitted to the relatively sparse data on air frosts at Earnscleugh. The actual fitting was carried out using the GLIM statistical package developed by Dr J.E. Nelder at Rothamstead (see Baker and Nelder (1978) for details). A plot of the results on normal probability paper is shown in Fig.3. The straight line shows the linear function

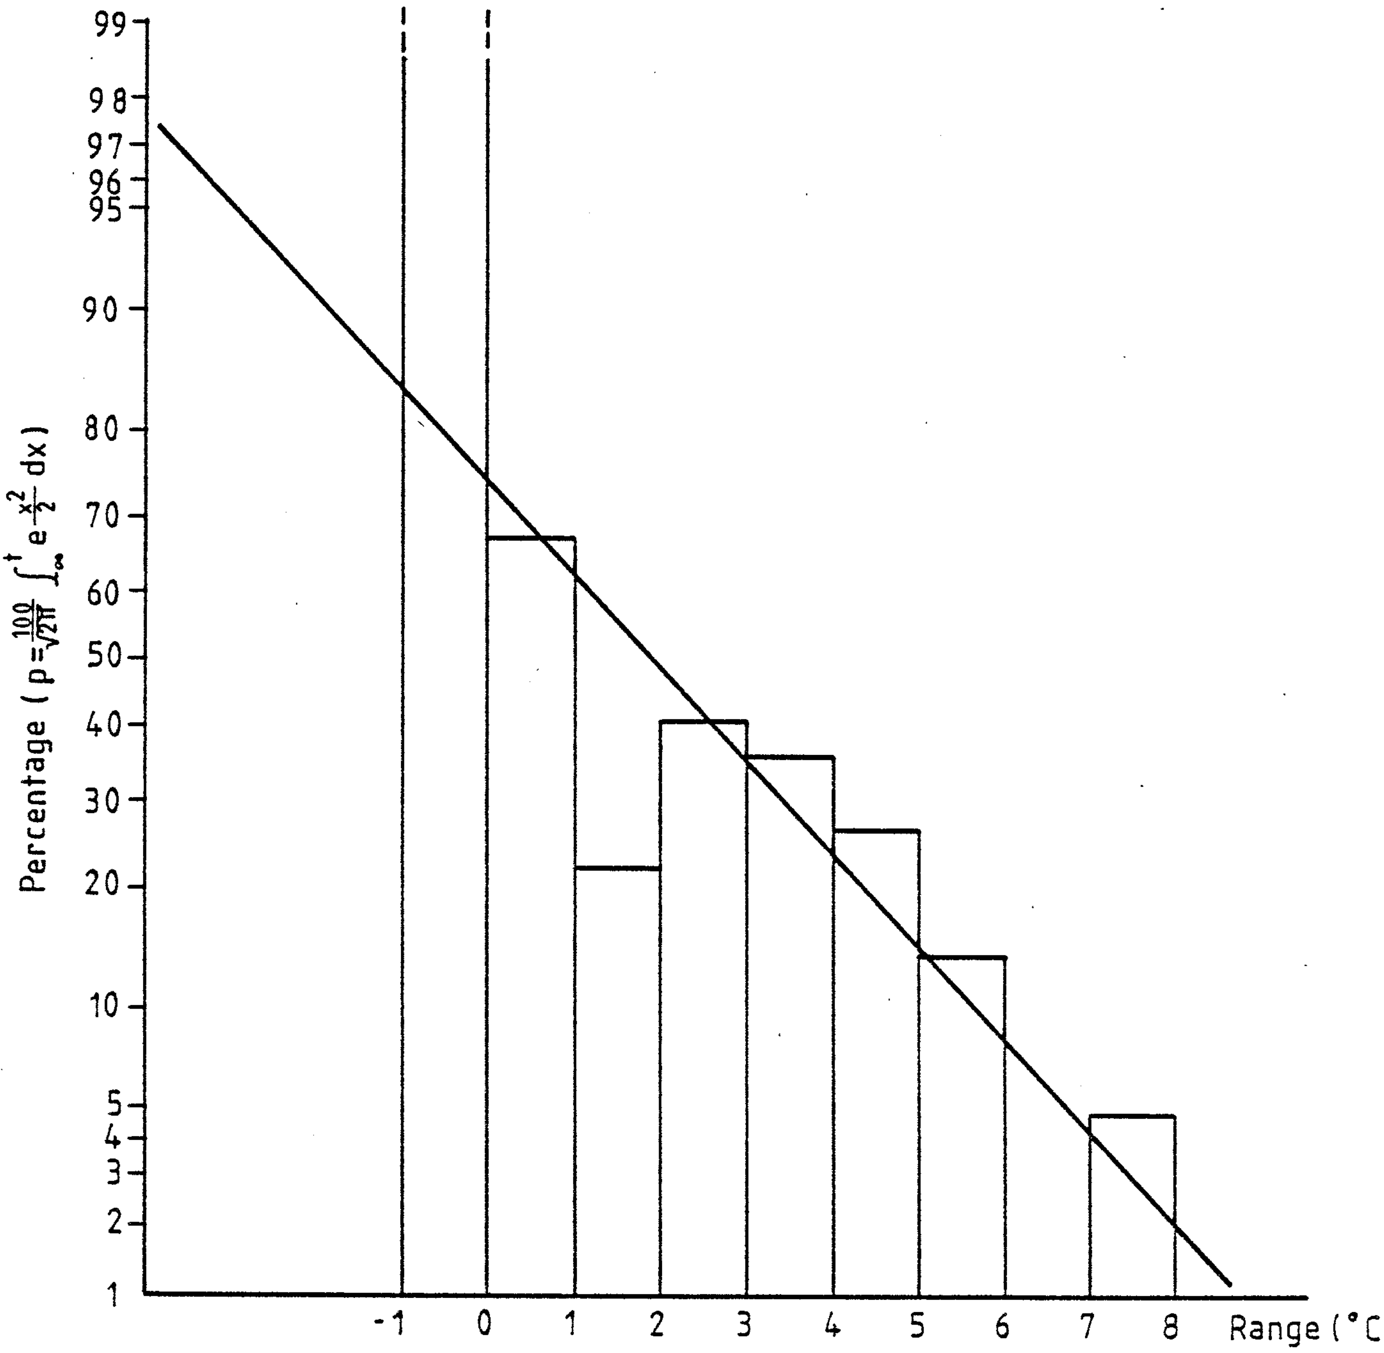

Figure 3. Actual versus Expected Probit Fit: all seasons. The figure shows the number of actual frosts at Earnscleugh as a proportion of the number of occurrences when the predicted overnight minimum temperature at Alexandra fell within a given 1°C temperature range (see Table 3).

obtained from the Probit fit, while the histogram bars represent the proportions of frosts observed between the indicated ranges for the predicted overnight temperature $\hat{T}_A$. The data (from the period 1968-71) is summarized in Table 3. Table 4 shows the coefficients, with their standard deviations, for separate seasonal fits as well as the overall fit.

Table 3: Number of Frost Occurrences as a Function of Predicted Minimum Temperature.

	Temperature Ranges (C^0) of Predictions $\hat{T}_A$									
	-1→0	0→1	1→2	2→3	3→4	4→5	5→6	6→7	7→8	8→
Number of Predictions	2	6	9	22	33	26	29	33	20	38
Number of frosts	2	4	2	9	12	7	4	0	1	0
Ratio (%)	100	67	22	41	36	27	14	0	5	0

The table gives the actual number of air frosts at Earnscleugh as a proportion of the number of occasions when the predicted temperature at Alexandra fell within a given temperature range.

Table 4: Seasonal Variation of Fitted Probit Coefficients

	α	S.D. of α	β	S.D. of β
1968	0.690	0.465	0.281	0.100
1969	1.157	0.682	0.443	0.145
1970	2.528	1.170	0.963	0.376
1971	-0.495	0.524	0.138	0.108
All	0.644	0.257	0.335	0.057

The table gives seasonal values of the coefficients α, β in the equation: -

Prob (airfrost at Earnscleugh | forecast minimum temp. at Alexandra = $\hat{T}_A$) = $P\{\alpha-\beta\hat{T}_A\}$

where P is the standard normal distribution function.

The fit to a straight line is perhaps as good as could be expected with the amount of scatter already illustrated with the regression data; it certainly does not suggest the use of any more sophisticated equation. Some discussion of departures from linearity is given in Sections 4 and 5.

3.3 Implementation

To implement the above procedure, the regression and probit algorithms have been made up into a subroutine which can be operated from a console on the forecaster's desk. The forecaster is provided with synoptic maps, and the mid-afternoon temperature and cloud cover report from Alexandra. He decides on the synoptic category, and enters this, together with the temperature and cloud cover information, into the console. The subroutine then provides him with: -

- (i) probability of air frost at Earnscleugh;
- (ii) forecast minimum overnight temperature at Alexandra;
- (iii) a recommendation as to whether or not to broadcast a prediction of overnight air frost at Earnscleugh.

He is then free to make use of these results in formulating or modifying his final forecast. An example of the output as it appears on the V.D.U. is shown in Fig.4.

Copy of VDU Text

```
**** FORECASTING AIR FROST AT EARNSCLEUGH ****
FIRST, I NEED TO KNOW WHETHER THE MSL PRESSURE FIELD AT 40S 160E
ON THE MIDDAY 733 IS ANTICYCLONIC OR CYCLONIC (IF IT IS HARD TO
TELL, LOOK AT AREA BETWEEN THERE AND THE SOUTH OF THE SOUTH ISLAND
FOR GENERAL CURVATURE) IS IT ANTICYCLONIC (Y/N) : Y
NOW PLEASE TYPE IN THE 3 P.M. TEMPERATURE AND CLOUD COVER AT
ALEXANDRA (TEMP IN CELSIUS/COVER IN EIGHTS, ESTIMATED IF NECESSARY)
: 10 5
THE ESTIMATED MINIMUM TEMP AT ALEXANDRA IS 1.90C THAT GIVES A
PROBABILITY OF 56% FOR AN AIR FROST AT EARNSCLEUGH
** I RECOMMEND A *FROST* FORECAST (P IS 21% OR MORE) **
NOTE THAT WITH THE ABOVE CUT  OFF OF 21%, PERFORMANCE HAS BEEN
                              FORECAST
                           FROST   NO FROST
                 FROST       38        3
       ACTUAL
                 NO FROST    68      112
```

Figure 4. Copy of Text as appears on VDU for a particular day.

4. Evaluation

To compare the performance of the probit model with traditional "Yes/No" procedures, the final step has to be taken, as suggested above, of deciding on a probability value that can be used as a threshold in determining whether or not to broadcast a frost warning. For any particular threshold probability, a 2 x 2 table can be drawn up, as in Table 1, from which a variety of performance indices can be calculated. A commonly used index in meteorological work is the Hanssen-Kuiper skill score, defined as: -

$$I = \frac{n_{11}}{N_1} + \frac{n_{22}}{N_2} - 1,$$

where n_{ij} is the entry in the i-th row and the j-th column, $(i,j = 1,2)$ of the table and N_i is the row total for the i-th row. This index is sensitive to errors in the prediction of rare events, a feature which recommends its use in the present case, although it also leads to a certain instability in performance.

For large samples, the performance of the skill index may be analysed as follows. Suppose first that the probabilities can be forecast without error, and that the rate of occurrence of occasions on which a forecast of probability p is made is given by some function f(p), normalized so that

$$\int_0^1 f(p)\,dp = 1.$$

Since frosts will actually occur on a proportion p of those occasions on which a forecast of probability p is made, the overall rate of occurrence of frosts is

$$\mu = \int_0^1 pf(p)\,dp.$$

Let us suppose that the decision to forecast a frost is made when the probability reaches a threshold level p_o. Then the proportion of occurrences on which a frost forecast is correctly made will be

$$\int_{p_o}^1 pf(p)\,dp.$$

Similarly, the proportion of occurrences in which a no-frost forecast is correctly made will be

$$\int_0^{p_o} (1-p)f(p)\,dp.$$

Hence the skill score for this choice of p_o will be: -

$$I(p_o) = \frac{1}{\mu}\int_{p_o}^1 pf(p)\,dp + \frac{1}{1-\mu}\int_0^{p_o}(1-p)f(p)\,dp$$

The maximum value can be found by differentiating, which yields the equation: -

$$\frac{1}{\mu} p_o f(p_o) = \frac{1}{1-\mu} (1-p_o) f(p_o)$$

whence $p_o = \mu$. Thus the skill score is optimized by taking the threshold probability equal to the average probability of occurrence of frosts.

While this derivation assumes complete accuracy in determining the probability of a frost, a similar result for the expected skill score holds under the much weaker assumption that the probability estimate is unbiased. It should be noted that the occurrence of a single maximum for $I(p_o)$ depends only on the continuity and differentiability assumptions for $f(p)$, and not on the form of $f(p)$.

Values for the optimum threshold based on empirical data vary from season to season, generally taking higher values in the seasons for which the prediction method worked well (1969, 1970) and lower values in the seasons for which the method was less successful (1968, 1971). The data for all four seasons together give a curve for which there are two local maxima, one for a threshold probability of 0.12 and the other for a threshold probability of 0.27. The overall proportion of occasions with air frosts is 0.21. These results are illustrated in Fig. 5 and Table 5.

Table 5: Optimal Hanssen-Kuiper Skill Score I and Optimal Threshold Probability p_o.

	I	p_o
1968	0.44	0.10
1969	0.76	0.13
1970	0.86	0.26
1971	0.38	0.14
All	0.54	0.12

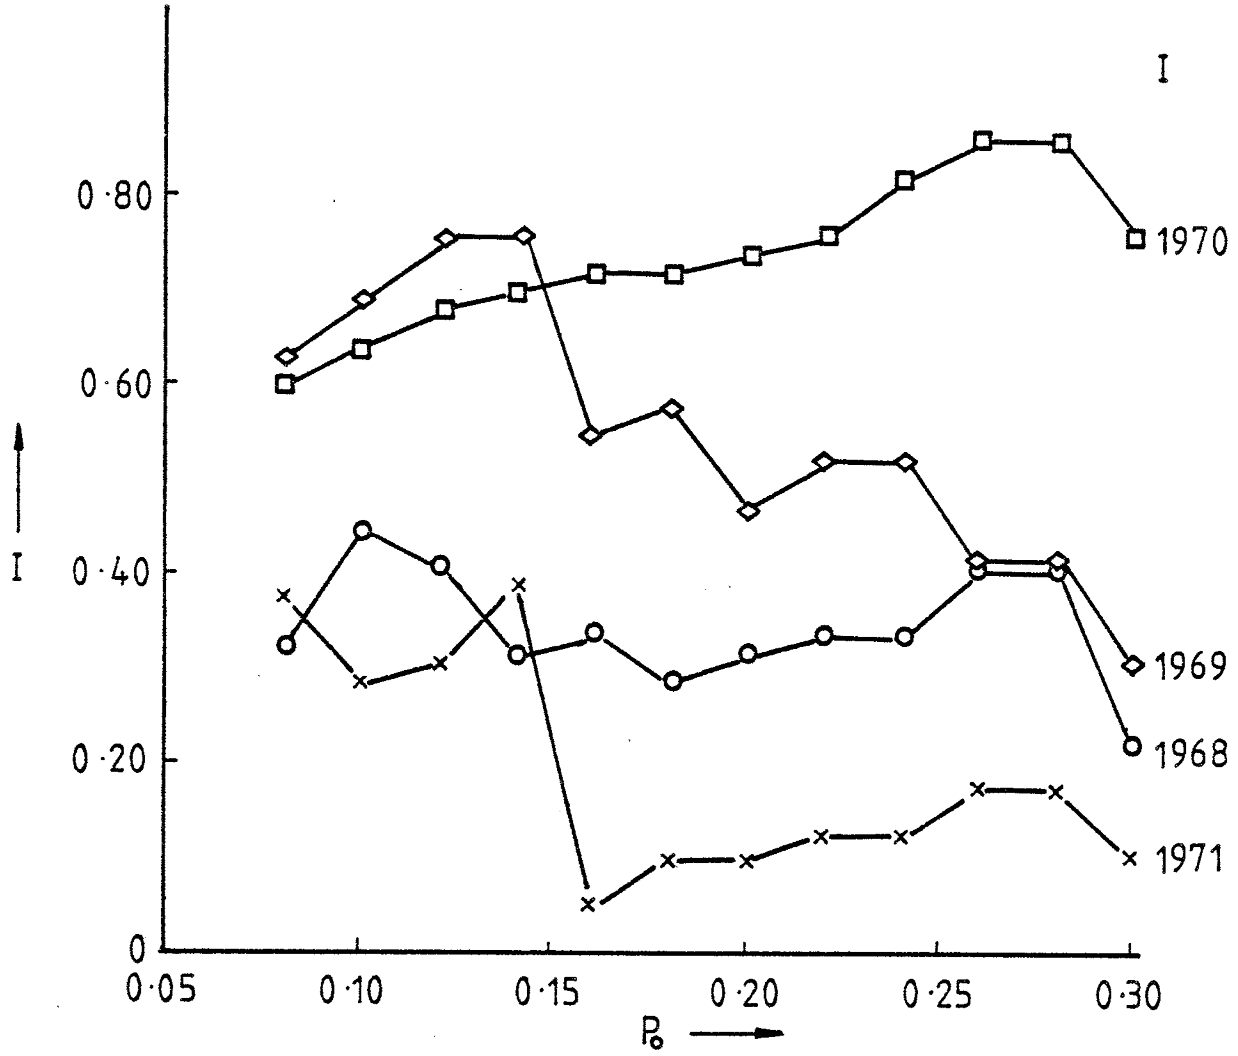

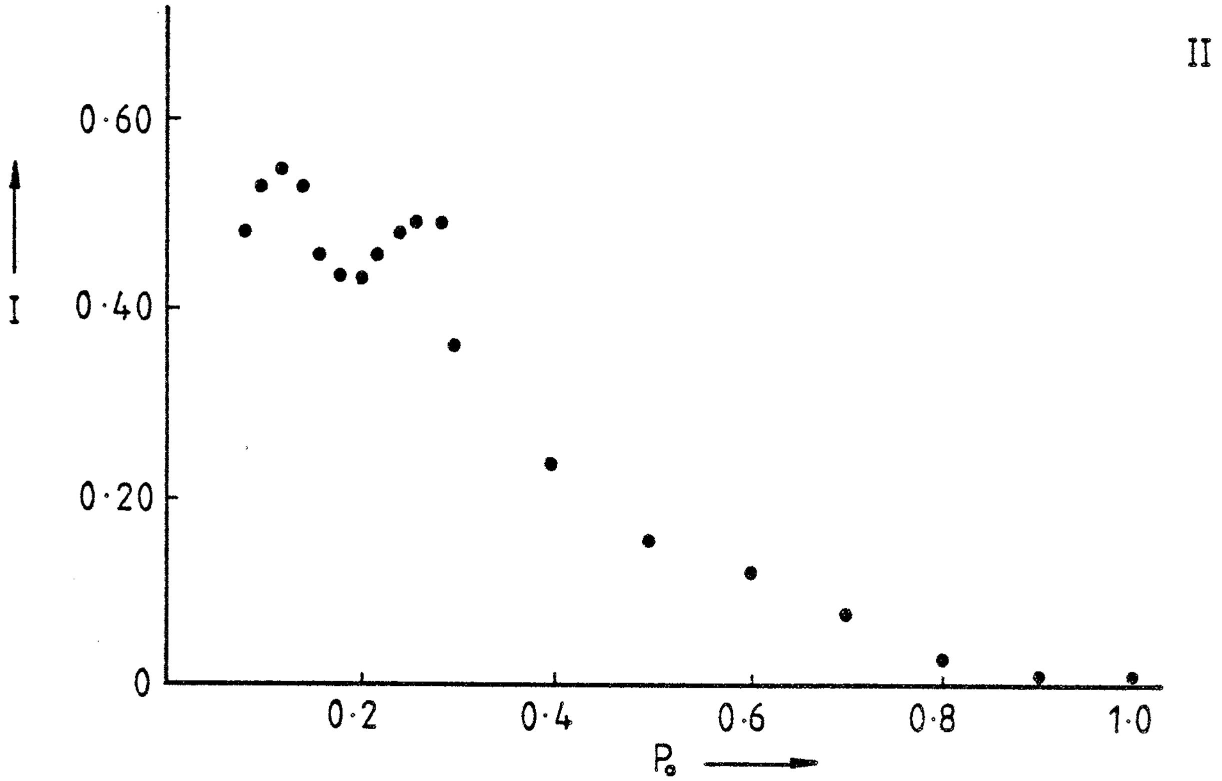

Figure 5. Hanssen Skill Score (I) as a Function of Probability Threshold (p_o) (empirical data)

I Variation with Threshold: individual seasons

II Variation with Threshold: seasons pooled.

For the purposes of implementation, a threshold of 0.21, close to the average probability of frost occurrence, was chosen in preference to any of the local maxima in Fig. 5. In fact the skill scores may vary only slowly in response to a change in threshold, so the optimal thresholds cannot be considered well-determined from the data, and the theoretical results obtained above may provide a more reliable guide than a threshold determined from scanty local data. Forecasters using the system have commented that even a threshold of 0.21 represents in practice a degree of overforecasting which they find hard to accept. This suggests that in practice factors over and above optimizing the skill score may have to be taken into account in determining a threshold level.

An alternative method of evaluating the performance of the model is to use a loss matrix which balances the cost of failing to forecast an air frost (and hence damaging the crop) against the cost of fighting frosts unnecessarily. Difficulties with this approach are that the costs are not known precisely and will in any case vary from orchard to orchard according to the type of equipment used, the nature of the crop, etc. Nevertheless a qualitative impression of the performance can be obtained by examining the expected cost under a loss matrix of the form:

		Action	No Action
Actual	Frost	1	r
	No Frost	1	0

Here the unit of cost is taken as the expected cost of deploying and operating frost fighting equipment for one night, and r measures in these units the expected cost of damage to the crop when a frost occurs and no action is taken.

Table 6 shows the variation of overall expected cost with threshold probability p_o and r. The entries have been calculated for the 1968-71 data according to the formula:

$$C = n_{11} + r.n_{12} + n_{21} = N_1 + (r-1)n_{12} + n_{21}$$

where n_{ij} has the same meaning as before. It will be seen from this table that the optimal thresholds p_o for the skill score are close to the values which minimize the expected loss over a broad range of values of r. The bimodality of the Hanssen-Kuiper index shown in Fig. 5 is naturally reflected in the structure of Table 6, with two valleys of minimum loss evident, corresponding to the local maxima of skill score. The actual departure from non-optimality, measured by the increase in expected loss for a threshold of 0.21 relative to the actual mimimum value of loss for a given r can be seen to be quite small for a range of values of r, being of the order of 10-12% for r = 1 to 5. For larger r the differences become quite significant (about 46% for r = 10) and a choice of p_O = 0.10 - 0.12 would be more appropriate. The table confirms the impression that 0.21 should be a suitable choice of threshold for implementation purposes for moderate values of r, but suggests it would be inadequate for severe potential loss r>> 5 say.

Table 6: Variation of Total Expected Loss, as a Function of Threshold Probability p_0 and Loss Measure r.

Threshold Probability	Loss Measure r							
	1	2	3	4	5	6	8	10
p_0 = 0.08	129	130	131	132	133	134	136	138
.10	116	118	120	122	124	126	130	134
.12	109	112	115	118	121	124	130	136
.14	104	109	114	119	124	129	139	149
.16	100	109	118	127	136	145	163	181
.18	94	105	116	127	138	149	171	193
.20	91	103	115	127	139	151	175	199
.22	86	98	110	122	134	146	170	194
.24	83	95	107	119	131	143	165	187
.26	75	88	101	114	127	140	166	192
.28	75	88	101	114	127	140	166	192
.30	69	89	109	129	149	169	209	249

The performance with 1978 and 1979 data is summarized in Table 7.

Table 7: Results of air frost predictions for 1978 and 1979

1978	Forecast		
	Frost	No Frost	
Frost	5	1	6
No Frost	27	37	64

I = .41

I = Hanssen-Kuiper index

1979	Forecast		
	Frost	No Frost	
Frost	4	0	4
No Frost	33	38	71

I = .53

5. Discussion and Conclusions

The results all suggest that the simple procedures outlined here represent too crude a simplification of the forecasting problem to be considered more than a first step in the evolution of a forecasting scheme. In this section we briefly review some of the departures from an adequate performance with a view to obtaining hints as to the reasons behind the unsatisfactory features and possible methods of correcting them.

One feature is the large variability from season to season. Clearly, the forecaster has a much easier task in some seasons than in others, no matter whether he is using objective or subjective criteria. This variation is evident in the parameters for the best-fitting probit lines as well as in the variations in optimal skill score and threshold probability. It is natural to ask whether these variations can be related to different weather patterns prevailing in the different seasons. Several indications of the general weather pattern were considered from this point of view, of which the most promising seemed to be the difference in average pressure anomalies between Christchurch (NE of the fruit growing region) and Invercargill (SW of the region). The difference gives an indication of the strength of the average E-W airflow over the period, higher pressure in Christchurch suggesting a dominance of anticyclones passing to the North of the region and causing a westerly airflow, higher pressures in Invercargill suggesting a dominance of anticyclones passing to the south of the region and causing a more easterly airflow. The relevant values are set out in Table 8, together with two other indications of general weather patterns, namely the number of occasions with pressure over 1015 millibars at Invercargill, and the number of occasions showing a strongly anticyclonic pattern over the South Island. Some examples of anomaly maps for different months are shown in Figure 6. The results confirm the impression that 1968 and 1971 (and to a lesser extent 1978) were characterized by more changeable, turbulent weather, while 1969 and 1970 were more settled.

Table 8: Indicators of Weather Pattern

	1968			1969			1970			1971		
	Sep	Oct	Nov	Sep	Oct	Nov	Sep	Oct	Nov	Sep	Oct	Nov
1. Anomaly Differences	0.4	1.7	0.6	0.4	-0.1	-0.4	0.4	0.7	-0.4	0.8	1.5	-0.4
2. Pressure Exceedances	4	3	1	5	16	12	5	7	11	3	2	5
3. Frequency of AC Situations	19	13	11	13	18	22	15	16	19	14	14	14

1. Tabulated values are the spatial differences Christchurch-Invercargill monthly pressure anomaly.

2. Tabulated values are the numbers of occasions in each month that the midday sea-level pressure at Invercargill exceeded 1015 millibars

3. Tabulated values are the number of occasions in each month that the analyzed sea level pressure map over the South Island was strongly anticyclonic.

Further insight into the behaviour of the model can be gained from an analysis of the departures from the straight line y = x in the plot of actual against forecast temperatures shown in Fig.2. The broad trends can be summarized as follows. The occasions when actual temperatures fell below forecast temperatures occurred mainly with lower temperatures, and were

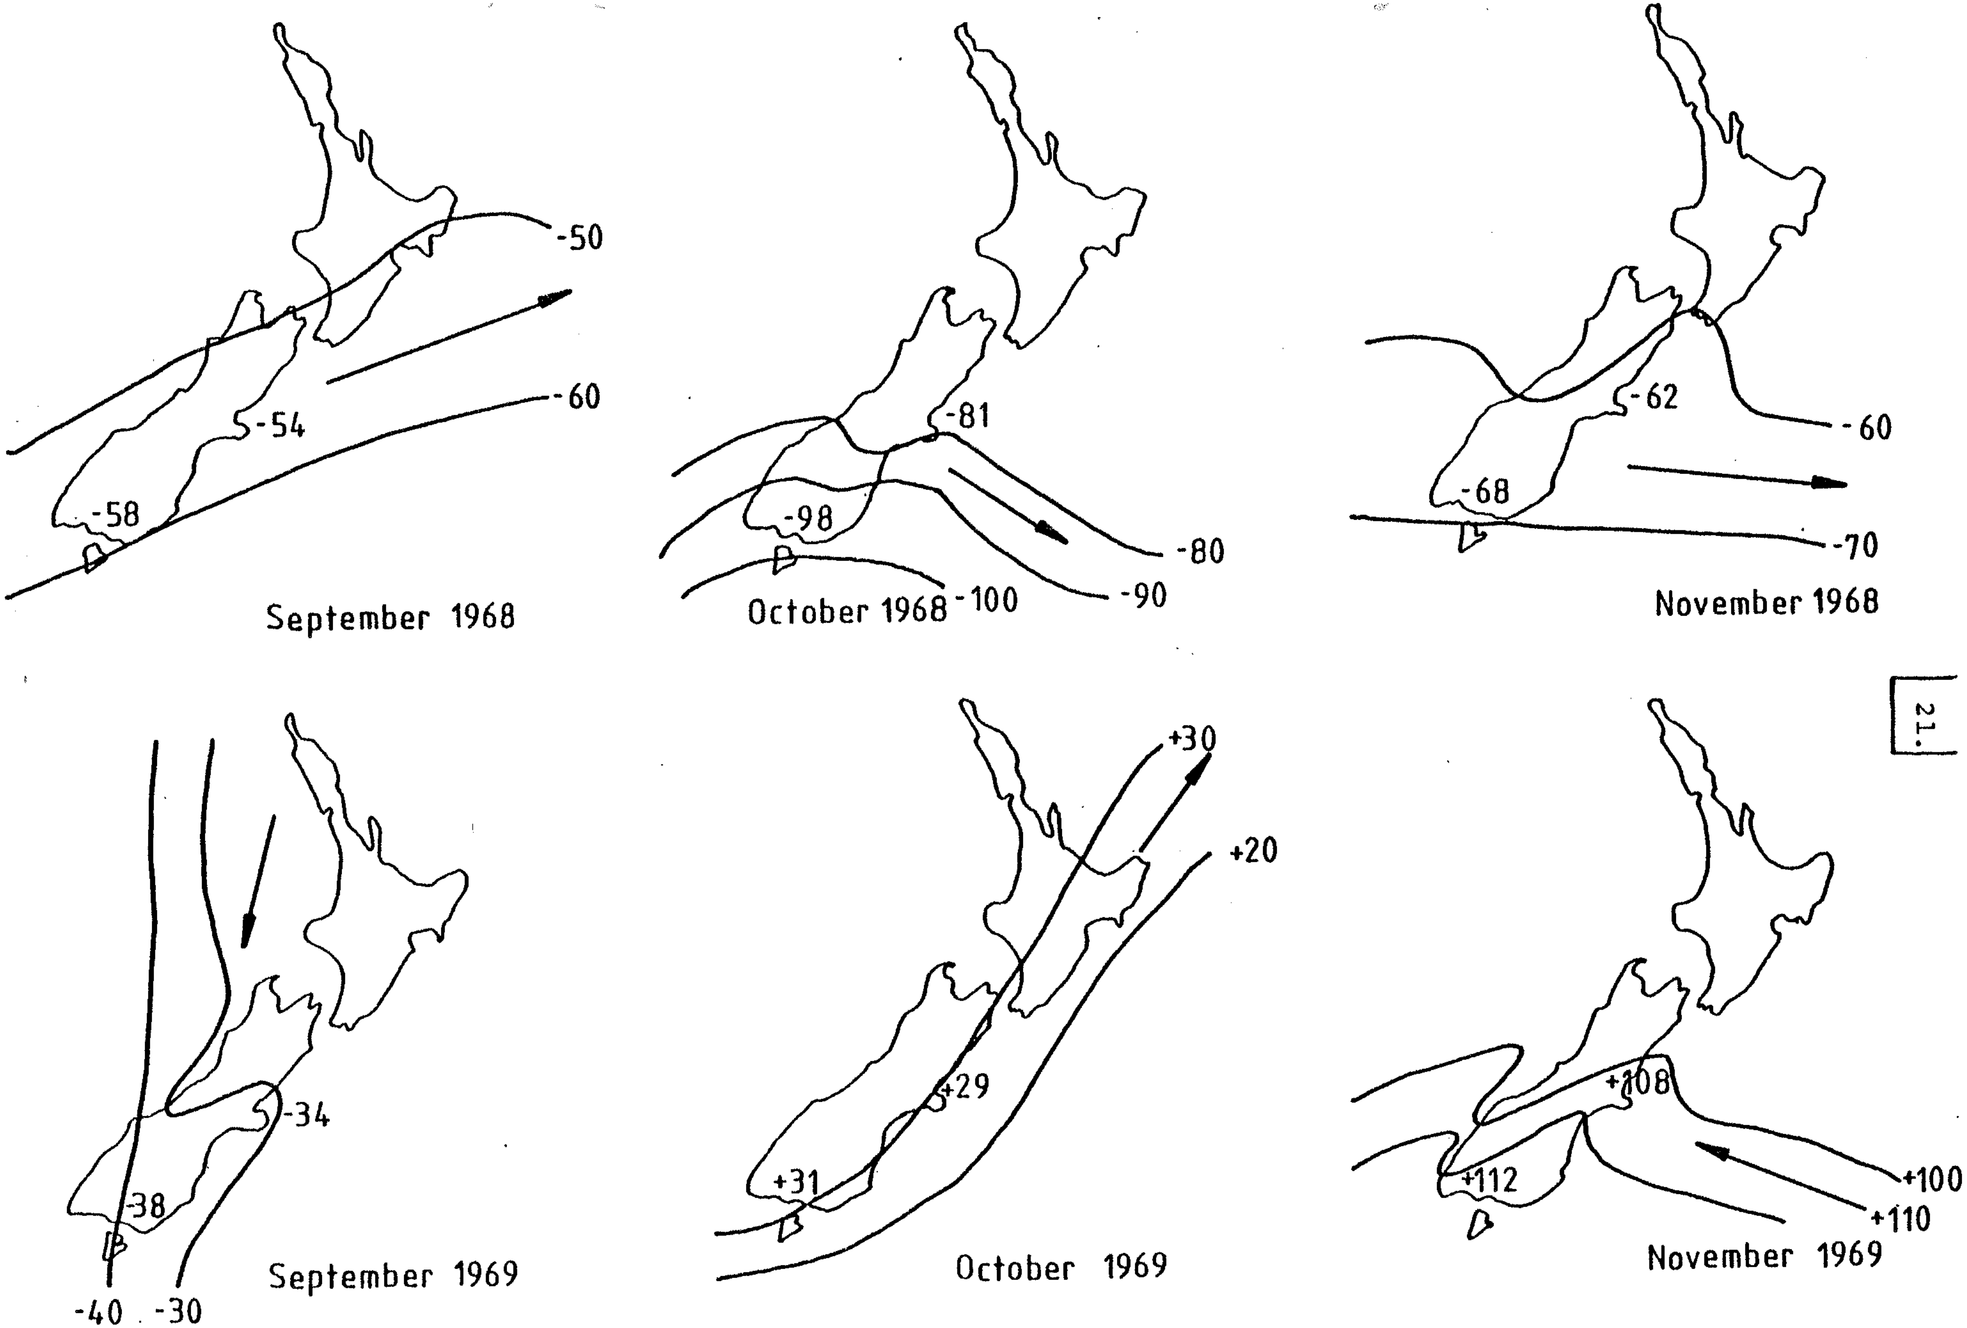

Figure 6. Examples of Pressure Anomaly Maps for years with good/bad performance by the Model.

Figures give the pressure anomalies (millibars x 10) at Invercargill and Christchurch, and those associated with smoothed contour lines.

Arrows indicate the direction of wind flow anomaly.

quite strongly associated with the presence of a southerly or southeasterly flow, with disturbances present. More specifically, a number of occasions on which Alexandra temperature forecasts were relatively high but frosts occurred at Earnscleugh were associated with the rapid onset of cold fronts from the southwest, in a manner not easily predictable from afternoon conditions. On the other hand, occasions when actual temperatures exceeded forecast temperatures occurred mainly in the higher temperature region, and were strongly associated with a northerly or northwesterly flow over the region, often with a trough approaching across the Tasman Sea. The presence of high cloud acting to reduce radiational cooling may also have been a factor here.

These remarks hint at the possibility of refining the analysis by the introduction of additional terms which would take into account dynamic effects of the types described above. The difficulty here is that the insertion of such terms would almost certainly require a more sophisticated classification of synoptic patterns, running thereby into one of the more difficult general problems of quantitative weather forecasting. For this reason, among others, we have not attempted to incorporate such refinements into the present study but have preferred to leave them as a subject for further research.

In summary, the conclusions of the present study may be set out as follows: -

(a) The probit model offers a simple, objective, but rather crude method of converting afternoon weather information into a quantitative estimate of the probability of an overnight air frost at a particular site.

(b) Although the skill scores from the method are modest, they are comparable with those for other forecasting situations of recognized difficulty (see, for example, Thompson (1968) who obtained values in the range 0.3 - 0.45 for 12-36 hours forecasts of rain in the main New Zealand centres). Certainly they represent an improvement over the existing procedure where air frosts are rarely mentioned.

(c) Together with a forecast of the expected minimum temperature the suggested procedures offer the grower both an indication of the likely severity of the frost and of the confidence he can place in the forecast.

(d) The main errors in the forecasts arise from failure to predict changes in the local weather conditions rather than from failure to predict the temperature fall under settled conditions. Substantial variations in the success of the method from season to season appear to be due to the relative frequency of unsettled conditions in the different seasons.

(e) Further research is needed to locate and correct the causes of incorrect forecasts by the present model and to increase its ability to discriminate effectively between frost and no-frost situations. In particular, further research is needed into methods of incorporating some dynamic terms into the algorithms for computing the probabilities of frost occurrence.

A more objective approach to the synoptic classification problem by using a principal component approach to the pressure field is now under investigation, using 1979/1980 data.

(f) There are some aspects of importance to the grower which have not been considered in this report but could perhaps be handled by a similar approach, for example the problem of forecasting the "first frost" in a series of frost episodes.

(g) Several techniques for the problem of handling dichotomous events in a probabilistic model formulation have been investigated, and the problem of objective determination of a suitable threshold value has been simply treated.

Acknowledgements

The authors would like to acknowledge the advice and support of the staff in the N.Z. Meteorological Service in the preparation of this paper. The suggestion of using the probit analysis facilities of GLIM for this problem came from Mr W. Armstrong. The paper is based on work undertaken by the first named author as part of the requirements of the Diploma in Operations Research and Statistics from Victoria University.

References

Baker, R.J. and J.A. Nelder, 1978: *The GLIM Manual.* (Copyright of Royal Statistical Society).

Finney, D.J., 1971: *Probit Analysis* 3rd Ed., Cambridge University Press.

Kidson, E., 1927: *Protection of orchards against frost.* New Zealand Journal of Science and Technology, Vol.9, No.4, 213-225.

Thompson, D.C., 1968: Skill scores achieved in N.W.F.C. Precipitation Forecasting. *N.Z. Met.S.Tech.Note* No.177.

Veitch, L.G., 1959: *Frost forecasting in the Barossa Valley, South Australia,* Australian Meteorological Magazine, June 1959.