PSI RESEARCH DISCUSSION PAPER 4

Early Identification of the Long-Term Unemployed

This paper is partly based on research commissioned by the British Government's Employment Service and the Department for Education and Employment. We thank Dr David Firth of Nuffield College, Oxford, for valuable statistical advice, Martin Range of the Social Studies Faculty, Oxford University, for expert programming, and Professor John Copas of Warwick University and Professor Roger Tarling of Surrey University for advice on prediction methods. The usual disclaimer applies.

PSI RESEARCH DISCUSSION PAPER 4

Early Identification of the Long-Term Unemployed

Clive Payne and Joan Payne

Policy Studies Institute

ISBN 0 85374 782 2

PSI Report No 874

psi

Policy Studies Institute
For further information contact
Publications Dept, PSI, 100 Park Village East, London NW1 3SR
Tel (020) 7468 0468 Fax (020) 7468 2211 Email pubs@psi.org.uk

PSI is a wholly owned subsidiary of the University of Westminster

Abstract

The paper describes a feasibility study carried out for the Employment Service and the Department for Education and Employment in Britain to identify those individuals who, when first registering as unemployed, are likely to be at high risk of becoming long-term unemployed. A simple prediction scheme is developed that can be easily administered by officials who deal with unemployment registration. The scheme is based on a logistic model for the odds of long-term employment which is calibrated and cross-validated on a national cohort study containing data on employment histories. The paper concludes with a discussion of the effects of the decision rules that are chosen and of the prediction errors that arise when the prediction scheme is applied in practice.

Introduction

A high level of long-term unemployment has been a persistent problem in the British economy. The fact that the rate of long-term unemployment tends to fall only slowly in response to increases in the demand for labour lends support to the view that long-term unemployed people are not effective competitors in the labour market (Layard, Nickell and Jackman, 1991). This lack of competitiveness could be explained by decay of skills, loss of motivation, stigmatisation, poverty and social isolation, leading to ever-diminishing chances of finding work as the duration of unemployment increases (Jackman and Layard, 1991); by a mismatch between unemployed people's skills and the requirements of new jobs created in the economy (Evans, 1993); by insider/outsider theory (Lindbeck and Snower, 1985 and 1994); or by the operation of the social security system in Great Britain (McLaughlin, Millar and Cooke, 1989; Nickell, 1993; Schmitt and Wadsworth, 1993).

During the 1980s and 1990s, the British government introduced a variety of active labour market policies aimed at increasing the labour market competitiveness of long-term unemployed people. These have included practical help with job search; training in job search techniques; training programmes to improve skills, qualifications and 'job-readiness'; opportunities for work experience both with employers and on special projects; direct wage subsidies to employers who take on long-term unemployed people; and job trials. In general, eligibility for these programmes has been confined to people who are already long-term unemployed, for the very good reason that most people who become unemployed find work within a few months with minimal assistance from public sources – to offer intensive help to everyone on first becoming unemployed would be prohibitively expensive and wasteful. However, the eligibility requirement of a certain length of unemployment means that by the time that they receive additional help, unemployed people may be already suffering from some of the concomitants of long-term unemployment that in turn may make the task of helping them back into work more difficult.

Thus the question has arisen of whether, when people first register as unemployed at local offices of the Employment Service, it is possible to identify individuals at high risk of remaining unemployed for a long time. This would allow high risk individuals to be given extra help immediately, instead of waiting until they are already long-term unemployed. This paper reports on a study carried out on behalf of the Employment Service and the Department for Education and Employment in Great Britain into the feasibility of developing a practical method to predict which individuals, among a group of people newly registered as unemployed, will become long-term unemployed. The study is preliminary rather than definitive, relying on secondary analysis of an existing large data set. However, suggestions are made for further developmental work.

Statistical methods for prediction

The problem of predicting whether or not a particular individual will become long-term unemployed is of course quite distinct from that of establishing the general risk factors for long-term unemployment. A very large number of studies have shown that a wide range of individual background factors are statistically associated with long-term unemployment, for example older age, low levels of education and skills, ill health and an unstable employment history (see the review by Pedersen and Westergard-Nielsen, 1993). However, such findings are about group effects. To take a hypothetical example: we know that poor educational qualifications increase the risk of becoming long-term unemployed, and we may be able to estimate that in a group of people newly registered as unemployed, 60 per cent of those with poor educational qualifications will find work fairly quickly and 40 per cent will become long-term unemployed. However, this does not tell us which *particular* individuals among those with poor qualifications should be offered extra help early on, as we do not know whether any given individual will turn out to belong to the 60 per cent majority that gets work quickly or to the 40 per cent minority that remains jobless. We need a prediction method to enable us to calculate the risk of becoming long-term unemployed for each individual, and this method needs to be reasonably simple and viable to implement.

We also need to know the accuracy of the prediction method: what proportion of individual predictions are we likely to get right? An estimate of accuracy can be obtained either by statistical techniques or by an empirical follow-up study of unemployed people. Ideally both approaches would be used: statistical techniques would be used to devise the best way of calculating individual risks, and the method would then be tested in practice. In the present study, where we have only secondary data, we rely on statistical methods alone.

Statistical scoring systems for risk assessment are used for a variety of purposes. Examples include calculating the risk of reconviction of prisoners being considered for release on parole (Copas and Marshall, 1998), assessing the risk of Sudden Infant Death Syndrome (Carpenter, 1983), and risk scoring systems for debt default widely used by commercial companies. Each of these examples is concerned with the risk that some event occurs within a certain time horizon, and in each we have data on relevant individuals in the past. This includes information on whether and when the event of interest has occurred, and information on risk factors on which the probability of the event may depend. Each of the applications cited above uses a different method of constructing risk assessment scores, ranging from orthodox statistical modelling to informal approaches geared more to intuitive simplicity than to statistical validity.

In the present study we draw heavily on the work of Copas and Tarling (1986) and Copas and Marshall (1998) on developing risk scoring methods, particularly as applied to the risk of reconviction. Their method estimates the probability that a prisoner with a given set of characteristics will re-offend within a specified period, and this information is used by Parole Boards in England and Wales as one of the factors taken into account in deciding whether to grant parole.

In our case we are concerned with men and women who have just started a spell of unemployment, and we wish to predict which individuals will become long-term unemployed. In practice, the length of unemployment needed for eligibility for additional help has varied both over time and with the programme in question; for the present exercise we use an arbitrary criterion of 12 months. Our procedure is as follows:

1. we develop a set of alternative statistical models for the probability of remaining unemployed for 12 months or more based on our data set;
2. we choose the model which statistical tests indicate will give the most accurate predictions on future data sets;
3. we use this model to calculate a predicted probability of long-term unemployment for each individual member of the sample;
4. we adopt a ‘decision rule’ – that is, we choose a level of probability to use as the cut-off point for taking action;
5. we evaluate the prediction method by computing the number of correct and incorrect predictions that it yields.

Data

Our study used the 1958 British Birth Cohort, also known as the National Child Development Study (NCDS). This is a continuing follow-up study of all children who were born in Great Britain in the course of a single week in March 1958, plus people born abroad in the same week who entered Britain before the age of 16. Follow-up surveys were carried out at ages 7, 11, 16, 23 and age 33 (1991). At each new survey, attempts were made to trace all cohort members, not just those who had taken part in the previous round. At age 33, personal interviews were obtained with well over 11,000 respondents, representing 74 per cent of the original birth cohort, less those known to have emigrated or died. Further details of the survey design can be found in Social Statistics Research Unit (1993).

At age 33, cohort members were asked to record their work histories since first leaving full-time education. For the present study, we defined a spell of unemployment as a period of time in which the cohort member was 'not in a job' (either full-time or part-time) because he or she was 'unemployed seeking work', and during which he or she was registered as unemployed with the Employment Service. Our analysis was restricted to the period between the age 23 and age 33 interviews, so that information collected at age 23 could be used to predict the risk of long-term unemployment. We excluded 17 per cent of respondents who had incomplete or inconsistent work history data for this period. We then selected those who had had at least one spell of unemployment that started more than twelve months before the age 33 interview, thus avoiding the problem of censored data, but after the interview at age 23 (for people who had had more than one unemployment spell that met these criteria, we modelled their first spell). This produced a sample of 747 respondents who met our data requirements.

Nevertheless, NCDS was attractive for our purpose because it is nationally representative, has a good response rate, and has collected a very wide range of data. This breadth, covering many aspects of physical, educational and social development, meant that a wide variety of possible risk factors for long-term unemployment could be examined. The main disadvantage of our data was that we could not use age or ethnicity as predictor variables, although both are implicated in long-term unemployment in the UK. Age was excluded because all cohort members were born in the same week, and ethnicity because the information was not of sufficiently high quality.

Methodology

The problem of risk assessment for the duration of unemployment can be expressed formally in the following form: we wish to estimate the probability that, for an individual described by predictor variables x=(x1, x2, x3,...), an event E (for example, leaving unemployment) will happen in a time period of length t:

$$p(x,t) = P(E \text{ occurs in time } t \backslash t,x)$$

Two types of analysis are possible, depending on how we treat time.

i) Models for fixed time. Here we choose a fixed value of t and estimate the probability that the event of interest will occur after t. Models for fixed time are essentially binary regression models in which we model how the response depends on the risk factors, or predictor variables. In our sample of people who had started a spell of unemployment, the response variable had value 1 if the person remained unemployed for more than 12 months, and 0 if he or she left unemployment earlier than this.

ii) Models for varying time. Here duration models are used to analyse how risk depends on both the predictor variables, x, and time, t. For the sample on which we calibrate the model, we need the actual times (possibly censored) at which the unemployment spell ended. This approach has not been used much for prediction because of its complexity, although Copas (1995) describes an application to predicting reconviction for potential parolees. The output from the duration model is a probability distribution for the length of time spent unemployed, conditional on the predictor variables. For each combination of their values, it gives the probability that a client will be unemployed for any particular length of time we wish to specify.

Models for varying time provide more information than models for fixed time but it is harder to apply the results in a practical context, so for our feasibility study we opted for method (i).

A range of statistical methods has been used for predicting a binary response. These include point scoring (Copas 1993), logistic regression, linear discriminant analysis, and various non-parametric and more compute-intensive methods such as the density estimation methods reviewed in Titterington et al (1981), the binary tree searching procedure of Breiman et al (1984) and neural network-based approaches. In practical applications in the social sciences the main methods used for prediction in this context have been point scoring and logistic regression. Both these methods have the virtue of simplicity. Copas (1993) claims that 'simple methods often work remarkably well in practice', and have the further advantage of simplicity in understanding and using the results.

We chose the logistic model for the prediction method, as Copas (1995) argues that the point scoring technique is not based on any explicit model for the way the probability of the event depends on the predictor variables; only with a statistical model is it possible to assess whether the model gives a good representation of the data. The standard statistical model for a binary outcome is the logistic, where we model the logarithm of the odds of the event occurring after the fixed value of t. Estimated probabilities from this model (or odds if preferred) can then be calculated.

In statistical modelling we generally choose the model which best fits our data. However, we need to distinguish between retrospective fit (how well the predictions fit the data used to construct the model) and prospective fit (how well the predictions will fit a future data set). For prediction, we need to maximise prospective fit, as retrospective fit will tend to give too optimistic a picture of the performance of the prediction method in practice. This is because the choice of predictor variables and the coefficients estimated for them reflect not only their relationships with the probability of long-term unemployment in the population from which the sample is drawn, but also the effects of sampling and measurement errors.

The difference between retrospective and prospective fit is called shrinkage. It can be substantial, especially when large numbers of predictor variables have been screened before a final model is selected. There are two main ways of dealing with the shrinkage problem, namely shrinkage correction and cross-validation.

With shrinkage correction, a model is fitted to the whole data set to estimate how much the calibration will be biased in future samples. The aim is then to adjust the fitted model so that the predictors are, in some approximate sense, likely to be correctly calibrated when used in practice. Copas (1993) describes a shrinkage correction factor for the logistic model which indicates the extent of shrinkage, and has developed shrinkage corrections that can be used to adjust the fitted probabilities of the outcome.

However, the alternative method of dealing with the shrinkage problem, cross validation, has been much used for developing prediction models, and this is the method that we adopted. We split our sample of 747 people who had started a spell of unemployment into two random halves, the 'construction' sample and the 'validation' sample. The construction sample was used to estimate the coefficients of the possible prediction models, but the choice between models was made by selecting the model which fitted the validation sample best, using the change in Deviance test (the log-likelihood ratio test). This method, of course, halved the size of the sample available to fit each model. The construction and validation samples were formed by a random stratified selection procedure involving three binary variables: qualifications (good or poor), region (below or above average unemployment) and sex. The construction sample had 373 members, 105 of whom remained unemployed for more than 12 months; the validation sample had 374 members, of whom 95 remained unemployed for more than 12 months.

Prediction model

The next step was to develop a prediction model for the probability of remaining unemployed for more than 12 months, using the construction sample. Initially we chose variables that other evidence (including more general models fitted to the full NCDS data set at an earlier stage of the work) suggested would be important predictors of unemployment length. However, we did not use variables on which local Employment Service offices have no information when someone first registers as unemployed, and which thus could not be used in a practical context. We also avoided choosing variables on which information could only be collected through possibly intrusive questioning by the claimant adviser, which clients might resent.

All the variables in our data set which satisfied these criteria were categorical. This gave an important practical advantage: we can represent the predictions as a table of predicted probabilities (or odds) of long-term unemployment for each combination of the chosen predictor variables. Thus the results are simple to use – the person applying the prediction method to real life cases has merely to refer to a look-up table in order to find the predicted probability for any particular combination of risk factors. If the model included continuous variables, some calculations would be needed at this stage, though a practical method has been developed for this (Copas and Marshall, 1998).

The predictor variables we used to develop the model included sex, plus a number of variables measured at age 23 and additional variables measured at the start of the unemployment spell. Those measured at age 23 included good educational qualifications (defined as at least five passes in national examinations at age 16), completion of an apprenticeship, receipt of formal job training, self-assessed state of health, whether buying a house, and region of residence. Those measured at the start of the unemployment spell included the reason for becoming unemployed, number of children, main source of income and type of job last held. We also used data on possession of a driving licence.

The initial logistic model that we fitted to the construction sample, using all the predictors listed above, had 22 parameters corresponding to the non-redundant categories of the predictors. The model gave a Deviance of 353.76 on 351 degrees of freedom; when fitted to the validation sample, the Deviance was 416.9.

However, given our small sample size, we needed to reduce the number of predictor variables in order to get big enough sample numbers in each cell of the cross-tabulation, defining each possible combination of risk factors to give a stable prediction for individuals identified by that combination of variables. This meant eliminating some predictor variables, even though they were statistically significant. Note that in doing this we were not seeking to find the most important causal or explanatory factors behind long-term unemployment. We were simply trying to identify the most efficient predictors of future long-term unemployment, prediction in the context of our feasibility study being a purely pragmatic exercise. The method we used was to fit a sequence of reduced models, removing some predictor variables altogether and combining selected categories for others. Decisions about changes to the model were made on the basis of significant reductions in the Deviance when the candidate model was fitted to the construction sample.

The model we finally chose, with five binary predictors, is given in Table 1. When fitted to the validation sample it had a Deviance of 389, which, being roughly equal to the degrees of freedom, indicated a good fit. The association of the first two predictor variables, poor educational qualifications and high regional unemployment, with long-term unemployment is very familiar and needs no further comment. The third, lack of a driving licence, has been found to reduce the chances of leaving unemployment in a number of recent studies using different data sets in the UK (for example, Payne et al, 1996; Payne et al, 1999). As well as being an indicator of poverty, lack of a driving licence restricts job search by limiting mobility, makes some workplaces difficult to access (especially for those living in rural areas or where the job involves unsocial hours when public transport is not available), and rules out a range of jobs of which driving is an essential part.

The model also shows that those whose main source of income at the start of their spell of unemployment was means tested benefits were more at risk of becoming long-term unemployed than others. There are a number of possible reasons for this. Those dependent on benefits were less likely than others to have resources (such as a telephone or smart clothes) that make job search easier. They were also less likely to have a partner or other household member in work, and so tended to have poorer links with work-based networks. In addition, their reservation wage was likely to be higher, as, if they took a low paid job, their income in work could in some cases be lower

Table 1: *Final reduced logistic regression model for long-term unemployment (construction sample)*

	Estimate	*t ratio*
Constant	0.09	
good educational qualifications	1.00	–
poor educational qualifications	1.85	1.84
regions with low unemployment	1.00	–
other regions of GB	2.62	2.99
driving licence	1.00	–
no driving licence	3.42	2.06
main source of income not means tested benefits	1.00	–
main source of income is means tested benefits	2.06	2.55
other reasons for unemployment	1.00	–
reason for unemployment is pregnancy	2.35	2.64
Number of observations	373	
Scaled deviance	404.93	
Degrees of freedom	367	

Note: Coefficients are reported in their exponentiated form, with the base or reference category set to 1.00.

than their income out of work. The final variable retained in the reduced model was becoming unemployed through pregnancy, which increased the risk of remaining unemployed for more than 12 months. This association (which of course applied only to women) probably involved a range of mechanisms, some of which would no longer hold today because of changes in legislation. However at the time when members of our sample became unemployed, receipt of unemployment benefit in the UK depended on a less rigorous test of whether someone was actively seeking work than is the case today, and it was not unusual for pregnant women who had given up their job and had no intention of returning to work nevertheless to claim unemployment benefit. In addition, during the earlier part of the period covered by our study, there was much weaker legal protection against dismissal on grounds of pregnancy and only very limited entitlement to maternity leave, and both these factors were likely to make pregnant women vulnerable to long-term unemployment.

Having chosen our model, we used the validation sample to compute the predicted probability of long-term unemployment, assuming the model to be correct, for each combination of values of the predictor variables. These are given in Table 2 for men and in Table 3 for women. Thus, for example, according to our chosen model, newly unemployed men with good educational qualifications and a driving licence, whose main source of income at the start of their unemployment spell was not means-tested benefits and who lived at age 23 in a region with low unemployment, had a predicted probability of remaining unemployed for more than 12 months of only 8 per cent. At the other extreme, newly unemployed men with poor educational qualifications and no driving licence, whose main source of income at the start of their unemployment spell was means tested benefits, and who did not live at age 23 in a region with low unemployment had a predicted probability of becoming long-term unemployed of 68 per cent.

Tables 2 and 3 also give the difference between the predicted probability and the observed probability for the validation sample, plus the sample number in the cell on which this is based. As would be expected, the biggest discrepancies occur where there are very small sample numbers, while in general cells with large numbers have small discrepancies. However many cells have very small sample numbers, particularly in the table for women, and so the probabilities are estimated with very poor precision.[1]

Converting model results into individual predictions

So far we have only made predictions at the aggregate level. We can predict that, say, one in five of people with a particular set of characteristics will become long-term unemployed, but we do not know which one of the five it will be. Thus we must specify a decision rule to convert probabilities into decisions about whether to take action or not. A standard rule is to use a fixed threshold as the cut-off point – for example, a predicted probability of 50 per cent or more. As we shall see, the level at which this cut-off point is set is important.

Having chosen a decision rule, the prediction method can be evaluated by comparing predicted outcomes with observed outcomes in the validation sample, and calculating the proportion of outcomes wrongly predicted. However it is more illuminating to consider two different types of error:

1 These figures can be used to construct a simple measure of the fit, the Percent Misclassified. If a cell has a predicted probability of 0.19 with an observed probability of 0.13 based on a cell size of 64, then the predicted number of people in this cell who will be long-term unemployed is 64 x 0.19 = 12. In fact 64 x 0.13 = 8 were actually long-term unemployed in this cell. Thus 12 - 8 = 4 of the 64 were wrongly predicted. The Percent Misclassified is the number of people misclassified summed over all the cells, expressed as a percentage of the total number of people in the validation sample. In our application its value is 4 per cent.

Table 2: *Predicted and observed probabilities of long-term unemployment in the validation sample: men*

	Regions of low unemployment				*Other regions*			
	Driving licence		*No driving licence*		*Driving licence*		*No driving licence*	
	Not means tested benefits	*Means tested benefits*	*Not means tested benefits*	*Means tested benefits*	*Not means tested benefits*	*Means tested benefits*	*Not means tested benefits*	*Means tested benefits*
Good qualifications:								
(a) Predicted % probability	8	18	16	31	19	36	33	54
(b) Difference between (a) & observed probability	–2	–22	16	31	7	7	7	54
N in cell in construction sample	23	6	4	1	73	13	10	2
N in cell in validation sample	29	5	2	1	64	17	7	2
Poor qualifications:								
(a) Predicted % probability	15	29	26	45	31	51	48	68
(b) Difference between (a) & observed probability	–4	–5	–74	–	1	1	–19	–32
Cell N construction sample	17	0	4	2	44	18	18	3
Cell N validation sample	21	3	2	0	51	14	18	3

1. false positives: when we predict that people will become long-term unemployed but they do not do so, and
2. false negatives: when we predict that people will not become long-term unemployed, but in fact they do.

These two types of error are shown in Table 4. They have different implications. A high rate of false positives leads to wasted resources in giving early

Table 3: *Predicted and observed probabilities of long-term unemployment in the validation sample: women*

	Regions of low unemployment				*Other regions*			
	Driving licence		*No driving licence*		*Driving licence*		*No driving licence*	
	Not means tested benefits	*Means tested benefits*	*Not means tested benefits*	*Means tested benefits*	*Not means tested benefits*	*Means tested benefits*	*Not means tested benefits*	*Means tested benefits*
Reason for unemployment is pregnancy								
Good qualifications								
(a) Predicted % probability	8	18	16	31	24	–	–	–
(b) Difference between (a) & observed probability	–1	18	–17	–	–1	–	–	–
N in cell in construction sample	14	1	2	2	2	0	0	0
N in cell in validation sample	21	1	3	0	4	0	0	0
Poor qualifications								
(a) Predicted % probability	19	36	33	54	45	–	63	–
(b) Difference between (a) & observed probability	12	36	33	54	45	–	63	–
Cell N construction sample	32	1	3	0	3	0	0	0
Cell N validation sample	27	5	3	1	3	0	1	0
Reason for unemployment is not pregnancy								
Good qualifications								
(a) Predicted % probability	15	29	26	–	–	–	55	–
(b) Difference between (a) & observed probability	–5	–	6	–	–	–	–	–
N in cell in construction sample	8	1	7	0	0	0	1	0

N in cell in validation sample	10	0	5	0	0	0	0	0
Poor qualifications								
(a) Predicted % probability	31	51	48	68	61	–	76	–
(b) Difference between (a) & observed probability	–15	26	10	35	–39	–	–24	–
Cell N construction sample	33	1	15	1	3	0	5	0
Cell N validation sample	26	4	13	3	4	0	1	0

help to people who would have found a job unaided, while a high rate of false negatives means we fail to give early help to people who may have benefited from it.

Table 5 shows the extent of these different types of error when we apply our prediction model, comparing decision rules that specify a 50 per cent, a 75 per cent and a 25 per cent cut-off point.

Table 6 translates these figures into measures of prediction performance. Taking a cut-off point of 50 per cent, the overall error rate based on our model was 25 per cent. This is similar to the rate regularly obtained in criminological prediction (Copas and Tarling, 1986). False negatives formed 23 per cent of all negative predictions, but false positives formed 48 per cent of all positive predictions. As a result, if we acted on the 50 per cent decision rule, nearly half of the people who were offered early assistance to avoid long-term unemployment would not need it. In fact we would have correctly identified less than one in five of the people who went on to become long-term unemployed, and failed to identify (and so failed to help) more than four in five of them.

Table 4: *Types of prediction error*

	Actual outcome 12 months later:	
	Not unemployed	*Unemployed*
Predicted outcome 12 months later:		
Not unemployed	true negatives	false negatives
Unemployed	false positives	true positives

Table 5: *Prediction errors*

Total percentages

	Actual outcome 12 months later: Not unemployed	Unemployed	Total
Predicted outcome 12 months later:			
50% cut-off point			
Not unemployed	70	21	91
unemployed	4	5	9
Total	**74**	**26**	**100**
75% cut-off point			
Not unemployed	75	25	100
unemployed	0	0.3	0.3
Total	**75**	**25**	**100**
25% cut-off point			
Not unemployed	43	7	50
unemployed	31	19	50
Total	**74**	**26**	**100**

Both the overall error rate and the proportions of false negatives and false positives are very sensitive to the choice of cut-off point. If we increase the cut-off point to 75 per cent, Table 6 shows that we substantially reduce wastage of resources, but at the cost of failing to pick up nearly everybody who may benefit from early help and so making the prediction effectively worthless. If we reduce it to 25 per cent, we succeed in identifying more of those who become long-term unemployed, but at the cost of a considerable increase in wasted resources.

Discussion

The purpose of our study was to investigate the feasibility of developing a method that staff in local offices of the Employment Service could use, when someone first signs on as unemployed, to decide whether that person has a high enough risk of becoming long-term unemployed to make it worthwhile giving him or her additional help at an early stage. We did not try to find the most important causal factors behind long-term unemployment, but simply those indicators, in the data we had available, that gave the most efficient prediction. We conclude that although the methodology to do this is available and can be applied to this purpose, in practice the pattern of errors that emerged presented a dilemma. In choosing our decision rule, there was a

Table 6: *Prediction performance*

	cut-off point		
	50 %	*75%*	*25%*
Error rate:			
Overall (% of all predictions that are wrong)	25	25	38
False negatives (% of all negative predictions that are wrong)	23	25	13
False positives (% of all positive predictions that are wrong)	48	0	63
Performance:			
% of actual long-term unemployed who are correctly identified and helped	18	1	74
% of others who are wrongly identified and helped	48	0	63

trade-off between maximising the chances of identifying the future long-term unemployed, and minimising the resources that are wasted by giving early help to new clients who do not need it. Anderson (1988) has formulated a strategy for choosing the optimal decision rule if costs can be attached to right and wrong decisions. Nevertheless, in our data the level of errors gave cause for concern, whatever the decision rule.

These problems may to some extent reflect the small size of our sample and our inability to use two important predictors of long-term unemployment in the UK, namely age and ethnicity. For the idea of early intervention to be taken further, more development work needs to be done on a data set that is free of these problems. The UK's quarterly Labour Force Survey (LFS) would be suitable for this purpose, as it has large sample sizes, collects a wide range of data and follows up respondents over four further quarterly interviews. The LFS would have another important advantage. Changes in the nature of the demand for labour following on changes in the structure of the economy alter the relative risk of long-term unemployment for different groups, and so any prediction model can only remain valid for a limited period of time. Because the LFS is a continuing survey, it could be used to modify the prediction model as the labour market changed.

Further development work should also involve dummy runs in selected local Employment Service offices to validate the scheme on administrative rather than survey data. This would allow the prediction method to be tested out in real life, rather than relying solely on statistical techniques. In addition, it would be valuable to run an experiment to compare the accuracy of a formalised prediction method with that of the informal and partly intuitive judgements of experienced staff dealing with newly unemployed people. If the latter were found to be more accurate, it might be possible to incorporate some of the cues on which they base their judgements into a formal prediction method.

There is also scope for investigation into the most effective way of using the type of prediction method described here in decision making. The similar scheme used by Parole Boards in deciding whether to grant parole provides only one criterion amongst several, and is never the sole basis of a decision. In other applications (for example, predicting whether someone is likely to default on a loan), the formalised prediction method is used to screen out applications which are clearly one side or the other of the borderline for acceptability, leaving borderline cases to be investigated in more detail. Debate on these matters does not only involve statistical arguments, but takes in many other issues of policy and principle.

References

Anderson, T W (1988) *An Introduction to Multivariate Statistical Analysis*. New York: Wiley

Breiman, L, J H Friedman, R A Olshen and C J Stone, (1984) *Classification and Regression Trees*. Monterey: Wadsworth

Carpenter, R G (1983) 'Scoring to Provide Risk-related Primary Health Care: Evaluation and Updating During Use'. *Journal of the Royal Statistical Society*, Series B 150: 225–65

Copas, J (1993) 'The Shrinkage of Point Scoring Methods'. *Journal of the Royal Statistical Society*, Series C 42: 315–31

Copas, J (1995) 'On Using Crime Statistics for Prediction' in M Walker (ed) *Interpreting Crime Statistics*. Oxford: Clarendon Press

Copas, J and P Marshall (1998) 'The Offender Group Reconviction Scale: A Statistical Reconviction Score for USE by Probation Officers'. *Journal of the Royal Statistical Society*, Series C 47: 159–71

Copas, J and R Tarling (1986) 'Some Methodological Issues in Making Predictions' in A Blumstein et al (eds) *Criminal Careers and Career Criminals*, Vol 2. Washington DC: National Academy Press

Evans, G W (1993) 'Sectoral Imbalance and Unemployment in the United Kingdom 1963–84'. Oxford Economic Papers 45: 440–456

Jackman, R and R Layard (1991) 'Does Long-term Unemployment Reduce a Person's Chance of a Job? A Time Series Test'. *Economica* 58: 107–121

Layard, R, S Nickell and R Jackman (1991) *Unemployment: Macroeconomic Performance and the Labour Market*. Oxford: Oxford University Press

Lindbeck, A and D J Snower (1985) 'Explanations of Unemployment'. *Oxford Review of Economic Policy* 1: 34–59

Lindbeck, A and D J Snower (1994) 'Patterns of Unemployment: an Insider-Outsider Analysis'. LSE Centre for Economic Policy Research Discussion Paper No, 960, London

McLaughlin, E, J Millar and K Cooke (1989) *Work and Welfare Benefits*. Aldershot: Avebury

Nickell, S (1993) 'Unemployment and the Benefit System London'. Employment Policy Institute Economic Report Vol 7, No 13

Payne, J, S Lissenburgh, M White and C Payne (1996) *Employment Training and Employment Action: An evaluation by the matched comparison method*. Department for Education and Employment Research Series No 74

Payne, J, C Payne, S Lissenburgh and M Range (1999) *Work-Based Training and Job Prospects for the Unemployed: An Evaluation of Training for Work*. Department for Education and Employment Research Report RR96

Pedersen, P J and N Westergard-Neilsen (1993) 'Unemployment: A Review of the evidence from panel data'. *OECD Economic Studies* 20: 65–114

Schmitt, J and J Wadsworth (1993) 'Unemployment Benefit Levels and Search Activity'. *Oxford Bulletin of Economics and Statistics* 55:1–24

Social Statistics Research Unit (1993) *NCDS5: A Brief Outline*. London: City University

Titterington, D M, G D Murray, L S Murray, D J Spiegelhalter, A M Skene, J D F Habbema, and G J Gelpke (1981) 'Comparison of discrimination techniques applied to a complex data set of head injured patients (with discussion)' *Journal of the Royal Statistical Society* Series A 144: 245–175