The Economic and Social Review, Vol. 23, No. 1, October, 1991, pp. 73-92

Trends in the Share of Long-Term Unemployment in Ireland*

RICHARD BREEN
and
PATRICK HONOHAN
The Economic and Social Research Institute, Dublin

I INTRODUCTION

Unemployment in Ireland has long been characterised by two features: a relatively high rate of unemployment, and a relatively high proportion of the unemployed who are long-term unemployed (defined as unemployed for one year or more). The share of long-term unemployment in total male unemployment in Ireland, 1966-1990, is plotted in Figure 1. This shows several periods of rapid growth of which the increase during the 1980s was the greatest. Similarly, over the same period, the male unemployment rate, as a share of the labour force, has shown a considerable increase, most marked between 1979 and 1986.

There are a number of possible explanations of why long-term employment should have increased as a percentage of total unemployment over the whole period and especially during the 1980s. For example, it might be argued that, as unemployment increased, there was a widening of the difference between the likelihood of leaving unemployment having been unemployed for a short

Paper presented at the Fifth Annual Conference of the Irish Economics Association.

*We are grateful to Frank Barry and Paul Walsh for helpful comments as well as to participants at an ESRI Seminar.

or medium length of time, and the likelihood of having been long-term unemployed. This might have arisen because, faced with a larger pool of short-term unemployed job seekers, employers reduced even further their hiring of the long-term unemployed. Alternatively, one might argue that an increase in level of unemployment assistance payments relative to take home pay acted as a growing disincentive to seeking or accepting employment on the part of the long-term unemployed.

Figure 1: *Share of Long-term in Total Male Unemployment, 1966-1990*

In this paper we argue for a much simpler explanation of the variation in long-term unemployment over time, which does not involve any assumptions about behavioural change. We show that, for the most part, variation in the share of long-term unemployment in total unemployment can be explained by changes in the volume of newcomers to unemployment. Clearly, an increase in the inflow will at first reduce the share of long-term unemployment, but subsequently the share will return to the same steady state level if the conditional probabilities or hazard rates of escaping from short-term and long-term unemployment remain unchanged.

This paper shows that it is possible to explain most of the variation in the share of long-term unemployment during the 1980s on the basis of constant hazard rates. In particular the sharp increase in the share from one-in-three to about one-in-two between early 1983 and late 1985 is largely if not wholly attributable to the fact that the numbers of short-term unemployed had stopped growing in early 1983, and indeed began to decline from then on.

Taking the longer period 1972-1990, there has been only one sharp change in these hazard rates and this coincides with a change in the data series at the end of 1979.

Reference may be made to two related papers. Hughes and Walsh (1983) consider incentive effects and economic activity as determinants of unemployment survival probabilities. However their concern is with survival from durations of less than three months. The recent paper by Lehmann and Walsh (1990), which examines the impact of employment schemes on escape probabilities, looks also at the 6-12 month duration; unlike their study, the present paper does not attempt to throw light on the determinants of short-run variations in escape parameters.

The layout of the paper is as follows. We begin by discussing our data and notation and we set out the simple mathematics of long-term unemployment. We then describe our attempts to model the change in long-term unemployment in Ireland in the recent past and present our results. We discuss some issues arising from these results and conclude by drawing out the implications of our findings for policy and for forecasting the future level of long-term unemployment in Ireland.

II METHODOLOGY

2.1 *Notation and Modelling*

We use data on the numbers falling into three duration categories: unemployed for less than 26 weeks, unemployed for between 26 and 52 weeks and unemployed for more than 52 weeks. We call these the newcomers (or short-term unemployed), the medium-term and the long-term unemployed, respectively. The data used in this paper refers to males only.

This data has been collected twice-yearly, in April and October, since April 1980. Before then, data was collected on a quarterly basis in February, May, August and November. This quarterly series was initiated in 1966, and continued until the end of 1979. Other changes in the classifications and procedures for these statistics were introduced at the time of the reduction in frequency, as part of the general overhaul of unemployment statistics initiated by the Interdepartmental Working Group. Among the definitional changes in the series were the exclusion of all those unemployed over 65 years of age, and of those receiving unemployment payments by virtue of working systematically on short-time. We understand that one of the deficiencies identified in the old (pre-1980) procedures was that some reporting offices tended to mis-classify as newcomers, persons who became recipients of unemployment *assistance* having just exhausted their entitlement to the higher rates of unemployment *benefit* (after 15 months of unemployment).[1] We selected the May and November observations to link with the data for the 1980s. There are two missing observations, (the first half of 1970 and the first half of 1979). Since April 1989, a new method of reporting the data extends the durations for which details are available. This series is discussed in Section 3.3 below.

Let the numbers in each of the three duration categories at time t be denoted s_t, m_t and x_t respectively. Defining:

$$a_t = 1 - \frac{m_{t+1}}{s_t} \tag{1}$$

we may interpret a_t as the probability of escaping from unemployment in six months conditional on being unemployed for less than six months. Analogous conditional escape probabilities, or *hazard rates* may be defined for the two longer durations; we denote these b_t and c_t. Unlike a_t, they cannot be directly calculated from a single six-monthly transition with the available data, but noting that:

$$b_t = 1 - \frac{x_{t+1} - (1 - c_t)x_t}{m_t} \tag{2}$$

if we assume that the hazard rates decline linearly with maturity, according to parameter ϕ, we may write,

$$c_t = b_t - \phi(a_t - b_t), \tag{3}$$

where $\phi >= 0$. Then, if ϕ is known, b_t (and hence c_t) can be calculated from:

$$b_t = \frac{1 - [x_{t+1} - (1 + \phi a_t)x_t]/m_t}{1 + (1 + \phi)x_t/m_t} \tag{4}$$

In general, a time series of a_t will exhibit variability, as will time series of b_t, and x_t, generated conditional on some particular value ϕ. Figure 2 displays these for $\phi = 1$. The main interest of this paper is in exploring the impli-

1. From what follows, it will be clear that such a misclassification would tend to increase long-term hazard rates, and slightly reduce the newcomers' hazard rates.

cations of assuming that these observed hazards a_t, b_t and c_t are derived from a data generating process in which the underlying parameters, α_t, β_t, and γ_t are constant over long periods, and the variation in observed hazards is attributable to a stationary noise.

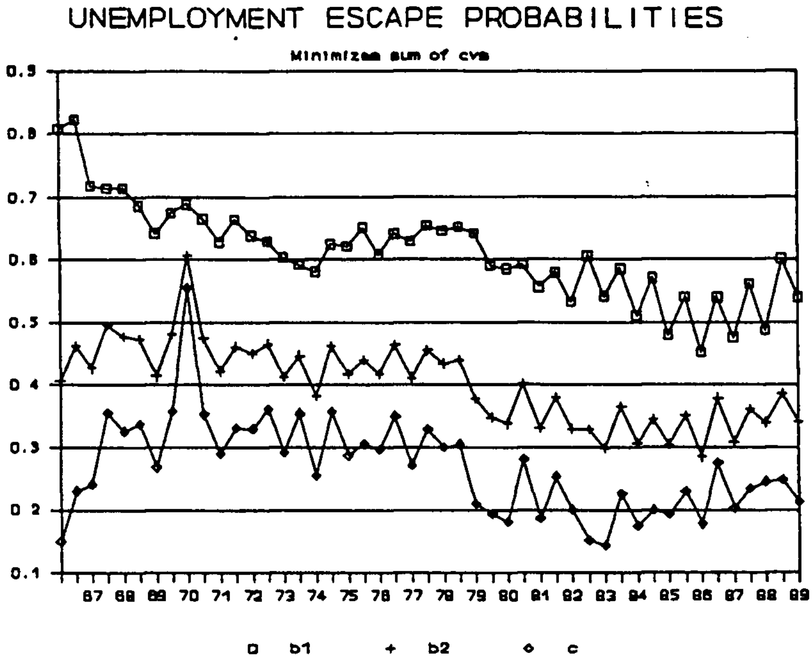

Figure 2: *Hazard Rates Assumed Declining Linearly after Six Months*

There are several ways of approaching the problem of estimating the underlying parameters and testing for their constancy. We examine two approaches here. Both are based on the model:

$$m_t = \alpha s_{t-1} + u_{1,t} \tag{5}$$

$$x_t = \beta m_{t-1} + \gamma x_{t-1} + u_{2,t}.$$

In one approach, the simulation approach, we carry out non-stochastic dynamic simulations of the model (5) using as inputs only the newcomers' data s_t. Because the actual data for neither m nor x is used, these are dynamic simulations. The constancy of the parameters is judged by reference to the goodness of fit of the dynamic simulation for x_t. In the other approach, the regression approach, the parameters of model (5) are estimated using actual data for all the variables. A variation of the regression approach substitutes out m, to obtain:

$$x_t = \alpha\beta s_{t-2} + \gamma x_{t-1} + v_t \tag{6}$$
$$v_t = u_{2,t} + \beta u_{1,t-1}$$

2.2 *Dynamics of Long-Term Unemployment*

To see how the share of long-term unemployment in total unemployment is related to the number of newcomers to unemployment we begin with a simple difference equation:

$$x_{t+1} = x_t(1 - c_t) + m_t(1 - b_t) \tag{7}$$

In words, the number long-term unemployed at time t+1 is given by the sum of the newcomers to long-term unemployment from the medium-term unemployed and the number of long-term unemployed at time t who did not leave unemployment over the period t to t+1.

Likewise, the number of medium-term unemployed at time t can be written as a function of the newcomers at t–1:

$$m_t = s_{t-1}(1 - a_{t-1}) \tag{8}$$

Combining (7) and (8) yields:

$$x_{t+1} = x_t(1 - c_t) + s_{t-1}(1 - a_{t-1})(1 - b_t) \tag{9}$$

If we assume that the coefficients a_t, b_t and c_t are constant over time, then, for an arbitrary starting date, t = 1, we can apply (9) iteratively to yield:

$$x_t = x_1(1 - c)^{t-1} + (1 - a)(1 - b)\Sigma_{k=1}^{t-1} s_k (1 - c)^{t-k-1} \tag{10}$$

for all t > 1. If the newcomers, s_t, is also assumed constant over time, (10) becomes:

$$x_t = \frac{(1 - a)(1 - b)}{c}(1 - (1 - c)^{t-1})s + x_1(1 - c)^{t-1} \tag{11}$$

and if we take the limit of (11) as T tends to infinity we get (using a bar to denote steady state values):

$$\bar{x} = \frac{(1 - a)(1 - b)}{c}\bar{s} \tag{12}$$

It is straightforward to generalise these results to situations in which the total unemployed are aggregated more finely according to duration (for example, into three month rather than six month intervals).

The number long-term unemployed has an equilibrium value which depends upon the newcomers to unemployment and the three escape probabilities. Equation (12) shows that the rate of change in the equilibrium number long-term unemployed with respect to a change in the newcomers is given by the ratio of the probability of remaining unemployed for one year, divided by the probability of leaving unemployment after one year. This ratio will typically exceed unity in situations where the escape probabilities are all small and decline with increasing duration of unemployment.

Finally, we can derive an expression for the equilibrium share of long-term unemployment in total unemployment. It is convenient, in this instance, to measure this share as the odds of long-term unemployment rather than the percentage. Thus, we write:

$$\frac{\bar{x}}{\bar{s}+\bar{m}} = \frac{(1-a)(1-b)}{(2-a)c} \qquad (13)$$

III THE EVIDENCE

Figure 1 shows the share of long-term unemployment in the total 1966-90. Superimposed on a generally upward trend are three major dips. The dips have their lowest points in 1971, 1975 and 1980. The numbers of long-term unemployed are plotted in Figure 3; it is noteworthy that although this series too has dips, they are not as pronounced as the dips in the share series, except at the very end of the series — since 1986 — when the numbers have fallen sharply.

Also on Figure 3 are the numbers unemployed for less than six months — the newcomers. Although we have also labelled this series the "inflow", it differs from the usual concept of inflow in that, because of short-spells, it is much lower than the number of persons who became unemployed during that six-month period. Furthermore, it refers to a six-month period, whereas inflow data in the international literature usually relates to much shorter periods. These differences should be borne in mind in any comparison of the present study with international analyses of inflow and outflow. The newcomers series has a different character to the others. It displays a small peak at 1971, and a larger peak at 1975 before returning in each case to a level of about 30,000. A third peak at 1982 is followed by a steady decline (apart from a seasonal sawtooth) until the end of the series. Though this decline is more gradual than those following the previous peaks, it is not clear that a new and higher plateau has been reached.

Figure 3: *Male Unemployment: Total and Less than Six Months*

For the 1970s, the juxtaposition of the newcomers and the long-term share suggests a relationship between the dips in the long-term share and the peaks in the numbers in the newcomers. For the 1980s, such a link is not so evident, but we will show that here too changes in the newcomers have been a crucial factor in influencing the growth and subsequent stabilisation of the long-term share. Specifically, the rapid growth in the number of newcomers in the early 1980s prevented the share of long-term unemployed reaching its steady state value; but as soon as the inflow of newcomers stabilised and then began to decline, the share of long-term unemployed rapidly jumped close to its steady state value.

Returning to Figure 2, showing the estimated hazards for each six-monthly transition (and on the assumption of a hazard rate declining with duration, $\phi = 1$), several points are apparent. First, the earliest four observations display a pattern quite different to the others in that the hazards move in opposite directions. They also provide by far the lowest escape probabilities from long-term unemployment and the highest from short term. We suspect that these observations may have been affected by teething problems with the collection of this data. Second, there is a sharp peak in the hazards computed for 1970. The CSO reported that long-term unemployment figures for this period were affected by the operation of Employment Period Orders. These first two comments lead to the conclusion that the early data is unreliable and analysis should begin about 1972. Third, there is a pronounced seasonal pattern, which seems to intensify in the 1980s. Fourth, hazard rates in the 1970s are generally higher than in the 1980s. Our main focus of interest is to discover whether it is possible to model long-term unemployment using a simple specification of the hazards, e.g., a constant, or a step function of time.

3.1 *Simulation Approach*

The simplest simulation model assumes that hazard rate parameters β and γ were constant over the whole period 1972-90. Feeding this into the model of the previous section and taking the average hazard rates as the fixed rate throughout generates the dynamic prediction plotted in Figure 4. Clearly this

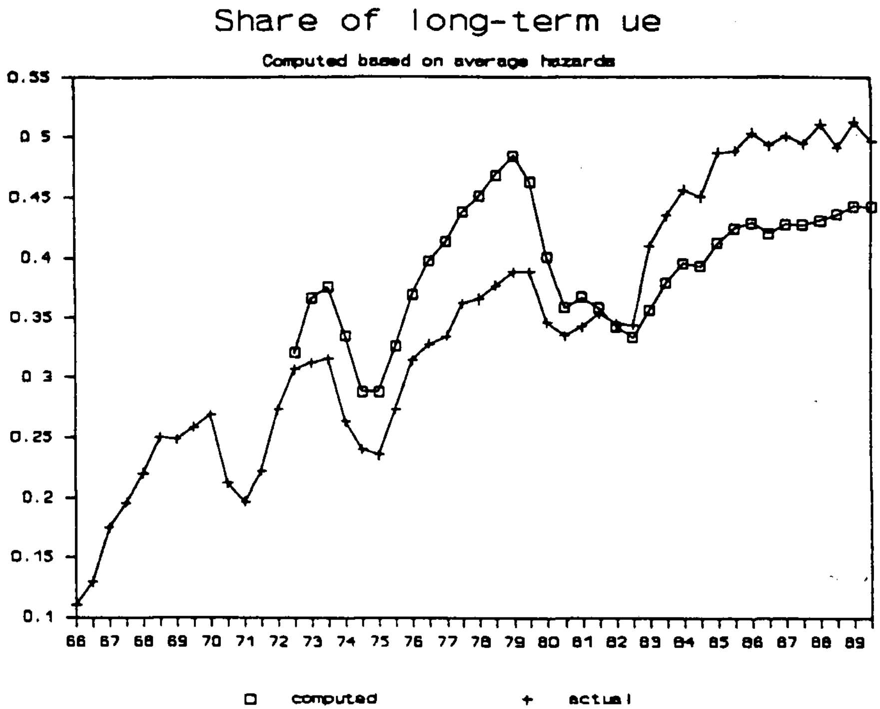

Figure 4: *Simulated Share of Long-term Unemployment using Average Hazard Rates*

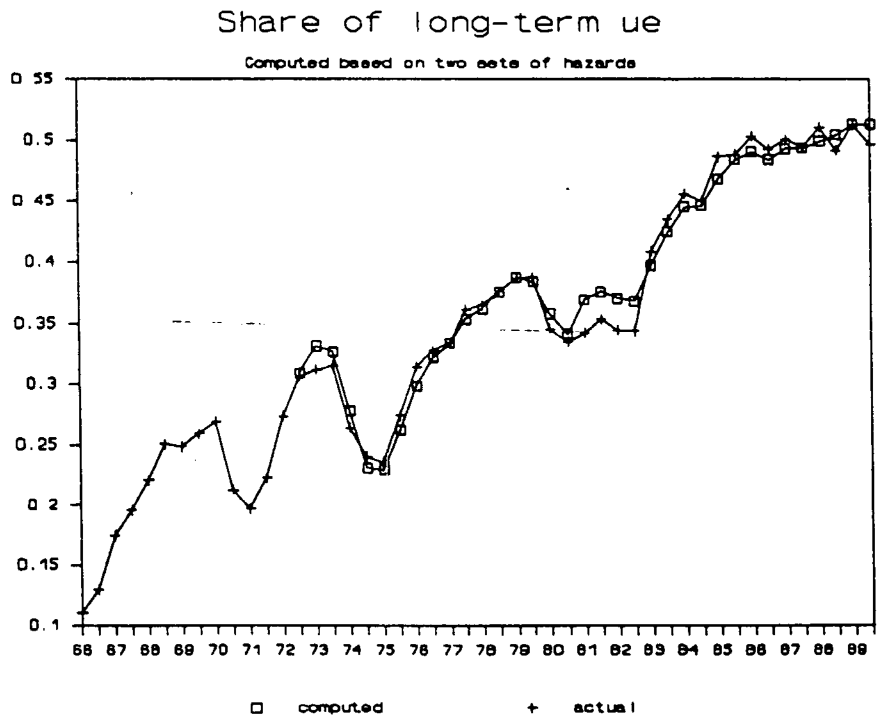

Figure 5: *Simulated Share of Long-term Unemployment using Two Sets of Hazard Rates, Broken at End-1979*

is not a very good fit: it is not satisfactory to assume that hazard rates were constant throughout the 1970s and 1980s. Drawing on the previous section, we move to the hypothesis that the hazard rates were constant except for a once-for-all jump at or about the turn of the decade. That is, the hypothesis that the parameters in each decade were constant. Using the decadal averages gives the dynamic prediction of Figure 5. This seems remarkably close. This figure was plotted for $\phi = 1$, but other values of ϕ give very similar results.

This simple approach appears to bring out the main features of the data. In practice, a slightly more complex iterative procedure was adopted. First, an initial value for ϕ was chosen as the value that minimises the sum of the variances of b_t and c_t for the 1980s. Using this value it was verified that the mean-squared error of the simulation was minimised with a parameter change after the second half of 1979. Using this split, a revised estimate of ϕ was obtained by minimising the mean-squared error of the simulation: this led to a value of $\phi = 1.1$ or 1.3 depending on whether the standard error of x or of its share in total unemployment was used. The precision of this estimate was low; for instance $\phi = 1$ gave almost identical standard error. Finally, we checked that, conditional on $\phi = 1$, the standard error was minimised at the date previously identified. While the variation of the simulation standard errors (Table 2) shows that the date of parameter change was estimated fairly precisely, the same cannot be said of ϕ.

Table 1: *Simulation Approach*

Dates	*Alpha*	*Beta*	*Gamma*	*Eq. shr.*	*SEE*	*RSQ*	*DW*
(a) Single set of parameters (phi=1)							
72h2-90h1	0.583	0.41	0.238	0.422	0.063	0.567	0.07
					14945	0.702	0.03
(b) Two sets of parameters (phi=1)							
72h2-792	0.63	0.465	0.3	0.325	0.013	0.982	0.95
80h1-90h1	0.551	0.374	0.197	0.496	3046	0.988	0.46
(c) Three sets of parameters (phi=1)							
72h2-79h2	0.63	0.465	0.3	0.325	0.012	0.984	0.96
80h1-82h2	0.59	0.394	0.198	0.47	2736	0.99	0.48
83h1-90h1	0.535	0.366	0.196	0.484			
(d) Three sets of parameters (phi=0)							
72h2-79h2	0.63	0.465	0.3	0.325	0.011	0.988	0.79
80h1-82h2	0.59	0.261	0.261	0.451	2022	0.995	1.25
83h1-90h1	0.535	0.243	0.243	0.497			

Notes: Eq. shr is the steady-state share of long-term unemployment for these parameter values.
SEE (RSQ, DW): Mean-squared error (respectively R-squared and Durbin-Watson statistic) for:
(upper figure) share of long-term unemployment;
(lower figure) level of long-term unemployment.

The estimated parameters are shown in Table 1 (these results are for the series cleaned up by excluding the over-65s before 1980; the raw data which includes the over-65s before 1980 but not after gives quite similar results; the figures are drawn for the raw data). Standard errors and R^2 statistics are computed both for the numbers of long-term unemployed and for their percentage share in the total. Durbin-Watson statistics are also shown as an indication of the serial properties of the residuals: note that a low Durbin-Watson does not necessarily signal biased parameter estimates here because of the method of estimation employed. The hazard rates for the newcomers (probability that they will escape unemployment in the next six months) falls from 0.630 in the 1970s to 0.551. The medium term hazard falls more sharply, from 0.465 to 0.374, while the long-term unemployed see their hazard falling from 0.300 to 0.197. The estimated steady-state share of long-term unemployed rises from 0.325 to 0.496, the latter figure being close to the figure at which this share actually peaked during the late 1980s.

Table 2: *Simulation Approach*

Mean Squared Errors for the Simulation					
For variations in date of parameter shift		*For variations in phi*		*For a second parameter shift*	
Date		*phi*		*Date*	
77h1	0.0244	0.0	0.0143	80h1	0.0142
77h2	0.0208	0.7	0.0129	80h2	0.0132
78h1	0.0188	0.8	0.0128	81h1	0.0135
78h2	0.0155	0.9	0.0128	81h2	0.0124
79h1	0.0134	*1.0	0.0128	82h1	0.0120
*79h2	0.0128	1.1	0.0128	82h2	0.0120
80h1	0.0132	1.2	0.0128	83h1	0.0118
80h2	0.0150	1.3	0.0128	83h2	0.0121
81h1	0.0176	1.4	0.0128	84h1	0.0134
81h2	0.0193	1.5	0.0128	84h2	0.0127
82h1	0.0222	2	0.0130	85h1	0.0140
82h2	0.0245	3	0.0133		
83h1	0.0272				
83h2	0.0295				

In each case the parameters α and β were taken as the within-decade sample means of a_t and b_t. A confidence interval was estimated around these values by assuming that these parameters were the true parameters and generating fifty sets of synthetic unemployment time-series, 1980-90, with these parameters, the actual s_t, and a disturbance term scaled to give the experienced goodness of fit. The parameter estimates that would have been derived from the synthetic data were plotted (Figure 6) allowing an approximate confidence interval to be drawn. This gives estimated standard errors of 0.006 and 0.002 for α and β.

Finally, we experimented with a second parameter shift during the 1980s. The best shift point seems to be after the second half of 1982; for $\phi = 1$ the improvement in fit is not dramatic, but we also find that $\phi = 1$ is no longer the optimal value, and a further improvement in fit is obtained by moving ϕ to the boundary of the permissible range at $\phi = 0$. The combination of a shift after 1982 and a zero value for ϕ is unattractive for several reasons. For one thing other evidence suggests $\phi > 0$ (see later). For another, it demands an implausibly sharp decline at the turn of the decade in the hazard rate for medium-term unemployed (from 0.364 to 0.261) compared with that for the newcomers (from 0.630 to 0.590). Taking into account also the relatively small reduction in the mean-squared error (from 0.128 to 0.106), we conclude that the simulations provide some evidence in favour of some small but systematic temporary change in the underlying data generating process in the early 1980s, but that a simple parameter shift does not seem adequate to capture this.

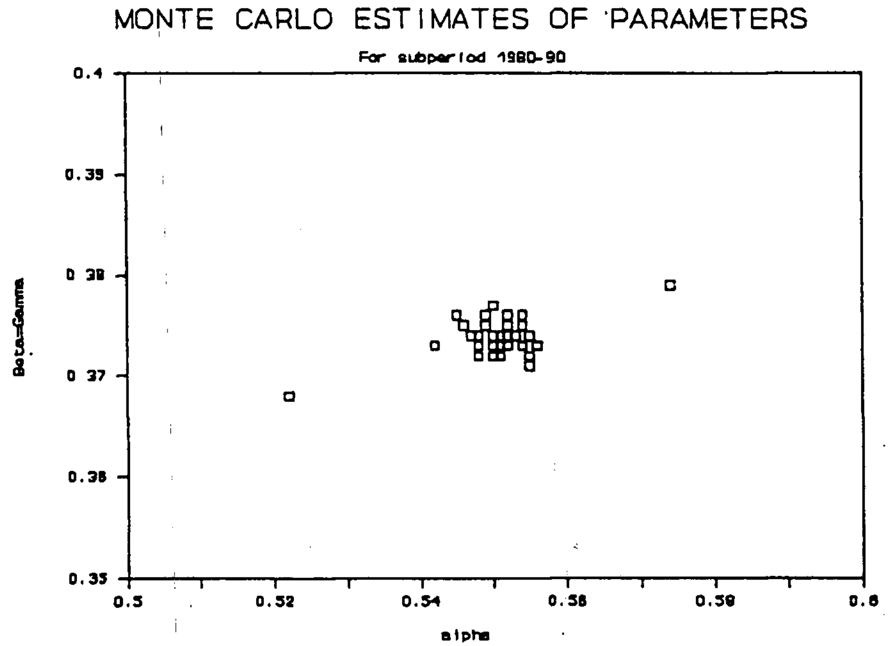

Figure 6: *Simulated Distribution of Parameter Value Estimates*

Our conclusion from the simulation approach is that the model (5) with just one shift in parameters, at the turn of the decade, allows us to simulate the evolution of the share of long-term unemployment remarkably closely, indeed with a mean squared error of just 1.28 per cent for the whole period 1972II-1990I. The timing of this parameter shift is fairly precisely determined at 1979-80.

Simple though it is conceptually, the simulation approach does not give us a ready access to standard hypothesis tests and inference results as does the more conventional regression approach to which we now turn.

3.2 *Regression Approach*

The regression equations estimated model (5) or (6), and tested for equality of the parameters at the break-points end-1979 and end-1982. That is to say, the parameters β and γ were modelled as the sum of a constant term, up to two shift dummies and a seasonal dummy. A selection of results are shown in Table 3 for both the structural model (5) and the reduced form (6).

The general findings were that the model gave very high R^2 values, but that the coefficients were not always estimated with precision. For Equation 5(a) the seasonal and 1980s dummies are estimated with precision and are robust to different specifications of the equation. Residual autocorrelation is a problem, and inclusion of a first-order autocorrelation parameter tends to make the 1980-83 dummy insignificant (regression B). (This autocorrelation is supposed to model unmeasured influences on unemployment.) The base value of α is about 0.61-0.66 depending on the season. This falls by about 0.09 at the turn of the decade. If there is a 1980-83 effect it is to defer some 0.03 of this fall.

Unconstrained estimates (not reported) of Equation 5(b) had rather low t-statistics for the various coefficients and it seemed appropriate to test equality constraints. Equality of the impact of seasonal and shift dummies on the parameters β and γ, were generally accepted. However, imposition of these equalities caused the estimated value of β to fall below that of γ, contrary to assumption (regressions D,E). Therefore the additional constraint of equating the base or constant values of β and γ was imposed (regressions F-J). (This constraint is not rejected conditional on the others being valid, but its imposition induces rejection of the equality on the 1980s dummy. However, allowing the 1980s dummy to take different values for β and γ produces implausible parameter values — regression G.) Equating β and γ is equivalent to setting $\phi = 0$, and a direct parallel can be made with the results of the simulation case where a second break was allowed. With all of these constraints imposed, and including a needed autocorrelation parameter makes the 1980-83 dummy insignificant (regression H). The base value of $\beta = \gamma$ is estimated in the constrained regression H at about 0.29-0.33 declining by about 0.11 at the turn of the decade.

Turning to the reduced form, Equation (6), once again equality constraints on the dummies were needed if high t-statistics were to be obtained. Since this equation does not allow separate identification of α and β, the problem of estimates of the latter being below those of γ did not arise. Here the sig-

Table 3: *Regression Results*

Structural Equations:

Equation 5(a): Dependent variable m

	coeff of	*s(–1)*			*rho*	*RSQ/DW*	*SEE/LL*
	base	*seas*	*D80*	*D80-83*			
A	0.340	0.054	0.096	–0.033		0.964	1402
	(30.6)	(5.5)	(8.2)	(2.4)		1.34	–301.1
B	0.338	0.054	0.092	–0.010	0.45	0.969	1330
	(19.6)	(8.3)	(4.7)	(0.5)	(2.6)	2.39	–298.7
C	0.338	0.053	0.089		0.51	0.968	1311
	(18.0)	(8.7)	(4.3)		(3.3)	2.47	–298.8

Equation 5(b): Dependent variable x

	coeff of	*m(–1)*			*x(–1)*		*rho*	*RSQ/DW*	*SEE/LL*
	base	*seas*	*D80*	*D80-83*	*base*	*D80*			
D	0.792	0.048=	0.139=	–0.046=	0.519	=		0.996	1845
	(8.5)	(5.5)	(6.2)	(2.2)	(8.7)			1.37	–310.2
E	0.842	0.050=	0.144=	–0.056=	0.494	=	0.36	0.996	1750
	(7.3)	(7.5)	(5.0)	(2.4)	(6.8)		(2.1)	2.07	–307.7
F	0.625=	0.043=	0.105=		=	=		0.995	1933
	(42.6)	(4.9)	(7.1)					1.43	–312.9
G	0.621=	0.052=	0.527	–0.064=	=	–0.567		0.997	1613
	(50.0)	(6.6)	(4.4)	(3.5)		(3.6)		1.64	–305.9
H	0.627=	0.045=	0.105=	–0.034=	=	=	0.36	0.996	1825
	(30.2)	(7.2)	(4.8)	(1.7)			(2.1)	2.07	–309.8
J	0.628=	0.044=	0.101=		=	=	0.29	0.996	1874
	(32.1)	(6.6)	(5.0)				(1.7)	2.02	–311.3

Reduced Form:

Equation (6): Dependent variable x

	coeff of	*s(–2)*			*x(–1)*		*rho*	*RSQ/DW*	*SEE/LL*
	base	seas	D80	D80-83	base	D80			
K	0.255	0.005=	0.120=	–0.030=	0.611	=		0.995	2088
	(5.8)	(0.1)	(5.6)	(1.7)	(9.8)			1.83	–314.5
L	0.206	0.010=	0.094=		0.685	=		0.994	2156
	(5.9)	(0.2)	(5.9)		(14.5)			1.82	–316.2
M	0.211	0.009=	0.096=		0.678	=	0.11	0.994	2179
	(5.5)	(0.2)	(5.3)		(13.0)		(0.6)	2.05	–316.0
N	0.258		0.120=	–0.030=	0.613	=		0.995	2054
	(7.0)		(5.7)	(1.8)	(10.3)			1.83	–314.5

Notes: These results relate to the data cleaned by removing all over-65s.
Estimation method: Ordinary least squares.
Rho: First order autocorrelation coefficient (estimated by ML).
"=" means that this coefficient is constrained to have the same value as the corresponding modifier of the coefficient of x_{t-1}.

Example: The fitted Equation D can be written:

$$x_t = (0.792 + 0.048\text{Seas} + 0.139D^{80} - 0.04D^{80\text{-}83})m_{t-1} + (0.519 + 0.048\text{Seas} + 0.139D^{80} - 0.046D^{80\text{-}83})x_{t-1}$$

where D^{80} is a shift dummy taking the value one from 1980 on, $D^{80\text{-}83}$ is a shift dummy taking the value 1 from 1980 to the end of 1982 and Seas is a seasonal shift dummy.

nificance of the 1980-83 dummy hinged on the inclusion of the seasonal dummy (regressions K,N). The estimated value of the product $\alpha\beta$ was about 0.27, and that of γ 0.40, both with a fall of 0.12 in the 1980s.

Overall, the regression results give somewhat less information than might appear at first sight. They confirm that there is a fall in the hazards between the 1970s and 1980s, but it remains unclear as to how this is distributed between the three key parameters α,β and γ. It is also not clear just what significance to attach to the 1980-83 experience. The multicollinearity with seasonal factors and with the autocorrelation correction clouds this issue. Only by imposing strong constraints (albeit ones that are not rejected by the data) can we discriminate between the various possibilities here. It becomes desirable to have some additional information. Fortunately, the new series with longer maturities begins to provide such information.

3.3 *The New Series*

The new series, for which three observations are available at the time of writing, allows further inferences to be made, especially about the degree to which hazards decline with maturity. That is because the "over one-year" category has been subdivided into three new categories: "1-2 years", "2-3 years" and "over 3 years". Denoting these categories x_1, x_2 and x_3, and noting that we have three observations, we may estimate the corresponding hazard rates γ_1, γ_2 and γ_3, by using the following implicit definitions.

This provides three equations with four unknowns, the three γ and β.

$$\begin{aligned} x_{1,t} &= (1-\beta)(1-\gamma_1)m_{t-2} + (1-\beta)m_{t-1} \\ x_{2,t} &= (1-\gamma_1)^2\, x_{1,t-2} \\ x_{3,t} &= (1-\gamma_2)^2\, x_{2,t-2} + (1-\gamma_3)^2\, x_{3,t-2} \end{aligned} \tag{14}$$

Adding the constraint $\gamma_2 = \gamma_3$ leads to estimated values $\beta = 0.471$, $\gamma_1 = 0.246$ and $\gamma_2 = 0.166$. These estimates are obtained with no degrees of freedom, but they do indicate a rather rapidly declining hazard rate with maturity. Indeed the estimate of β is rather higher than some of those estimated from the longer time series, and the γ's rather lower. As further observations become available, this or similar techniques can be adapted to allow relaxation of the constraint relating γ_2 and γ_3, but this constraint does not in fact affect the estimate of γ_1.

IV DISCUSSION

4.1 *Declining Hazards: Heterogeneity or a Causal Relationship?*

Declining hazard rates with maturity do not necessarily signify that any given unemployed person finds it more difficult to leave unemployment the longer he has been unemployed. The newly unemployed comprise persons with differing individual hazard rates; as time goes on, those with low escape hazards tend to remain unemployed while the others tend to escape, thereby lowering the observed group escape hazards. This composition effect is well known and has been widely studied in micro-data for other countries (cf. Kiefer, 1988). The composition effect is generally thought to be empirically important, but the existence of a distinct duration effect cannot be ruled out.

Recently, Jackman and Layard (1991) suggested a way of testing for the existence of a duration effect from aggregate data. Their method involves a lot of assumptions on whose validity in the Irish context we have no evidence. Nevertheless, it is worth reporting what their method yields on the data we have been using. In short, Jackman and Layard argue that, if there is no duration effect, group hazard rates for different durations should, in steady state, bear a constant proportion to one another. The trend in the Irish hazard rates has, however, been for a greater proportionate decline in the rates for longer maturities.[2] For example, the simulation estimates of α/γ increase from 2.1 in the 1970s to 2.8 in the late 1980s. If the Jackman and Layard approach is to be taken seriously, the fact that hazards decline with maturity cannot be wholly attributed to composition effects.

4.2 *What Happened at the Turn of the Decade?*

All the evidence points to a sharp drop in hazard rates at the end of 1979. What happened to cause the change? There is no shortage of possible explanations. First, aggregate demand factors could be relevant: end-1979 saw the beginning of a long recession with industrial employment falling for over eight years and unemployment rising for over seven years. The factors which induced these changes may well have influenced the hazard rates, though the 1970s also experienced a severe recession. Second, the general upward trend in replacement ratios for the long-term unemployed during the 1980s may have contributed, though this does not seem to have been sufficiently sudden to account for such a sharp break. Third, international factors may be adduced, including Ireland's membership of the EMS (March 1979) and more generally a worldwide shift to an anti-inflation stance can be dated to 1979

2. This is also what Layard and Jackman found for the UK. Incidentally, their data seems to show greater variation in hazard rates than the Irish data; it is noteworthy that their hazards also fall rather steeply around 1979-81, though this is not as pronounced a break in their data as it is in ours.

with, for example, the new Thatcher and Reagan governments and the flexible interest rate policy of the US Federal Reserve. Selecting among these possible interpretations will not be easy.

Our own view is that the change in series definition and in the procedures for collecting the data is likely to have been the most important source of the break. In particular, the timing coincidence is exact. Furthermore, the sharp change in seasonal patterns also points in the same direction. Although we have adjusted the earlier series by removing the over-65s, there were other changes, including the removal of those on systematic short-time working, and the change in the sampling months (though the latter can hardly explain the increase in the estimated equilibrium share of long-term unemployment, especially when there was no conspicuous immediate jump in the share of long-term unemployed when the sampling months changed). Above all, the procedural improvements may have had significant effects, as mentioned above.

V CONCLUDING REMARKS

Simple mathematics show that the number of long-term unemployed depends upon previous inflows to unemployment and on the conditional probabilities of escaping from unemployment. Increases in long-term unemploy-

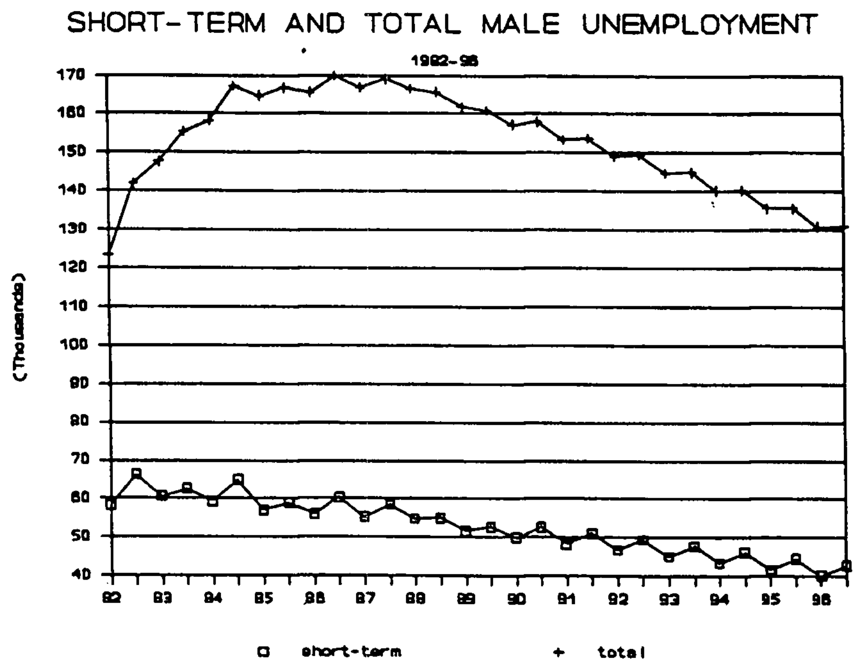

Figure 7: *Computed Long- and Short-term Unemployment to 1996*

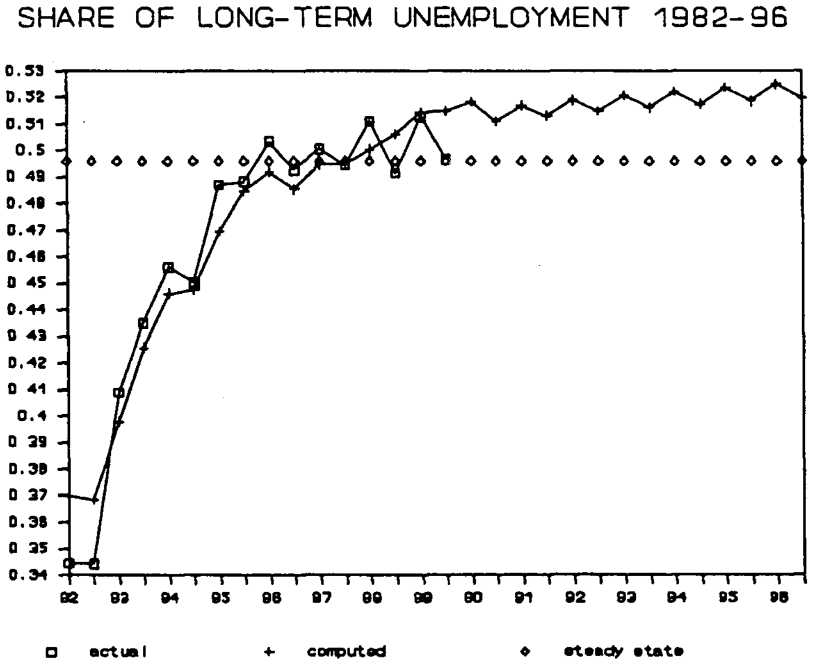

Figure 8: *Computed Share of Long-term Unemployment to 1996*

ment caused by behavioural changes — such as increasing discrimination against the long-term unemployed by employers or increasing disincentives to leave unemployment caused by changes in replacement ratios — should be reflected in changes over time in the conditional probability parameters. That we find these parameters to be broadly constant over long periods suggests that, at the aggregate level, neither the growth of long-term unemployment in the 1970s nor in the 1980s was due to such behavioural changes. Nor is there evidence of an impact of policy interventions of the 1980s on the hazard rates.

For each of these two periods variation in the inflow to unemployment is sufficient to account for the observed growth in the share of long-term unemployment.

The message of our model is therefore that long-term unemployment trends are heavily determined by stable hazard rates which allow a considerable degree of predictability. Indeed, if the number of newcomers is known, long-term unemployment can be predicted up to a 95 per cent confidence interval of +/– 4,000. As shown in Figure 7, the number of newcomers (unemployed for less than six months) has been falling slowly but steadily (apart from a seasonal factor) since early 1983. The growth in the long-term unemployed has simply been a manifestation of the low hazard rates. Since about 1988, the long-term share has reached its equilibrium steady-state value (Figure 8). A projection from the two-hazards simulation model is shown in Figures 7 and 8. If the number of newcomers were to continue to decline gradually, the long-term share would remain a little above its equilibrium value for the coming years, with the numbers of long-term unemployed declining gradually but steadily as they have done since 1988. By 1966 the numbers of long-term unemployed would be at 66,000 compared to their 1987 peak of 88,000.

REFERENCES

HUGHES, J.G., and B.M. WALSH, 1983. "Unemployment Duration, Aggregate Demand and Unemployment Insurance: A Study of Irish Live Register Survival Probabilities", *Economic and Social Review*, Vol. 14, No. 2, pp. 93-118.

JACKMAN, R., and R. LAYARD, 1991. "Does Long-Term Unemployment Reduce a Person's Chance of a Job? A Time Series Test", *Economica*, Vol. 58, No. 229, pp. 93-106.

KIEFER, N.M., 1988. "Economic Duration Data and Hazard Functions", *Journal of Economic Literature*, Vol. 26, No. 2, pp. 646-679.

LEHMANN, H., and P.W. WALSH, 1990. "Employment Schemes in Ireland: An Evaluation", *Economic and Social Review*, Vol. 23, No. 1, pp. 43-56.