Nordic Journal of Political Economy

Volume 26 2000 Pages 143-170

What Hides Behind the Rate of Unemployment? Micro Evidence from Norway

Knut Røed Tao Zhang

This article can be dowloaded from: http://www.nopecjournal.org/NOPEC_2000_a08.pdf

Other articles from the Nordic Journal of Political Economy can be found at:
http://www.nopecjournal.org

Knut Røed and Tao Zhang*

What Hides Behind the Rate of Unemployment? Micro Evidence from Norway

Building on a complete account of registered unemployment spells in Norway, we study how the composition of unemployment has developed over the last ten years. The total volume of unemployment has become more unequally distributed than before, but it is difficult to identify the 'losers' in terms of observed characteristics. There are no signs of low-education workers doing systematically worse. The most conspicuous change with respect to observed characteristics is that the relative outflow rates for older workers have deteriorated sharply. JEL-codes D39, J64.

For most European countries, the low rates of unemployment that prevailed until the first OPEC shock in 1974 are nowadays considered completely infeasible. Twenty years ago, unemployment in Europe stood at around five per cent of the labour force, and this level was regarded as devastatingly high. Today, only utopians seem to believe that Europe will ever return to unemployment rates much below this level. What has changed? Why does an achievement that after all was accomplished by most European countries thirty years ago now seem all but impossible? Numerous explanations have been offered in the literature (see e.g. Bean, 1994, for a recent survey). Some of these explanations are 'aggregate' in nature; i.e. they are embedded in the framework of a representative agent model. They typically build upon some sort of a wage bargaining- or efficiency wage model in order to identify the driving forces behind the rise in unemployment. The problem with the representative agent models is that they cannot explain the disproportionate distribution of unemployment spells, which is in fact one of the most conspicuous features of the European unemployment experience. Hence, much of the focus has turned towards explanations that are 'relative' in nature, i.e. explanations that not only seek to account for the level of unemployment, but also its distribution.

Some of these explanations build on the idea that macroeconomic shocks that cause transitory high unemployment may have

* The Frisch Centre for Economic Research, Oslo. We wish to thank the Ministry of Labour and Government Administration for financial support and Harald Goldstein, Oddbjørn Raaum, Steinar Strøm and Rolf Aaberge for helpful comments.

long-lasting structural effects in the labour market, either because the period of high unemployment entails 'sorting' mechanisms that systematically change the productivity distribution among unemployed workers (Johansen, 1982), and perhaps leave a fraction of the labour force 'stigmatised' (Vishwanath, 1989; Blanchard and Diamond, 1994); or because individual productivity and search effectiveness are adversely affected by longer periods of unemployment (Phelps, 1972; Hargreaves Heap, 1980; Pissarides, 1992). An alternative hypothesis, that has become popular in recent years, is that relative wages in Europe have failed to adjust to a more dispersed distribution of individual productivities (OECD, 1994; Krugman, 1994). Consequently, some low-skilled workers are entitled to wages that exceed their expected productive contribution. At the same time, the social security system provides a minimum living standard that for some workers is higher than what they would have achieved if they were to be paid according to their own productivity. The result is that some workers have become almost unemployable, and these consistently unemployed workers push up the aggregate rate of unemployment.

One implication of these 'relative' explanations is that today's aggregate rates of unemployment are not really comparable to the rates that prevailed some decades ago. The unemployment pool is no longer a pure 'reserve army' of readily employable labour (if it ever was), it is also (increasingly) a storehouse for workers who either do not compete for vacant jobs at all, or who do so with very low efficiency. Hence, the distribution of unemployment has become more unequal. Aggregate unemployment rates are more and more associated with relatively few workers being unemployed for relatively long periods of time.

Attempts to find out whether or not unemployment has become more unequally distributed among different skill groups are typically based on reported aggregate unemployment rates for workers with different educational attainment. The bottom line is that it has not. For example, Nickell and Bell (1995; 1996) find that the unemployment rates for workers with low education have *not* increased relatively more than the unemployment rates for workers with high education. They conclude that the rise in European unemployment is basically a skill-neutral phenomenon. However, there are a number of problems associated with the analysis of relative unemployment rates. First, apparently similar *rates* may conceal important differences with respect to incidence and duration. Second, it is difficult to identify the relevant measure of skills; in particular, it does not necessarily correspond to the crude educational measures used to calculate skill-specific unemployment rates. And third, it is not obvious how unemployment rates for different groups should be compared, when the overall unemployment rate has increased sharply. If for example the high-skill unemployment rate increases from one to two per cent, while the low-skill unemployment rate increases from six to 12 per cent, the ratio of the two unemployment rates is unchanged, while the difference between them has doubled.

The aim of this paper is to take a look behind the aggregate rates of unemployment, in order to throw some light on the contributions of 'aggregate' versus 'relative' unemployment theories. For this purpose, we have at our disposal register and survey data describing the more recent development of the microanatomy of unemployment in Norway. The Norwegian labour market possesses many of the structural characteristics often claimed to be responsible for high unemployment in Europe, such as generous

Figure 1.
The Norwegian UV-curve. Vacancy- and unemployment rates 1961–1998.

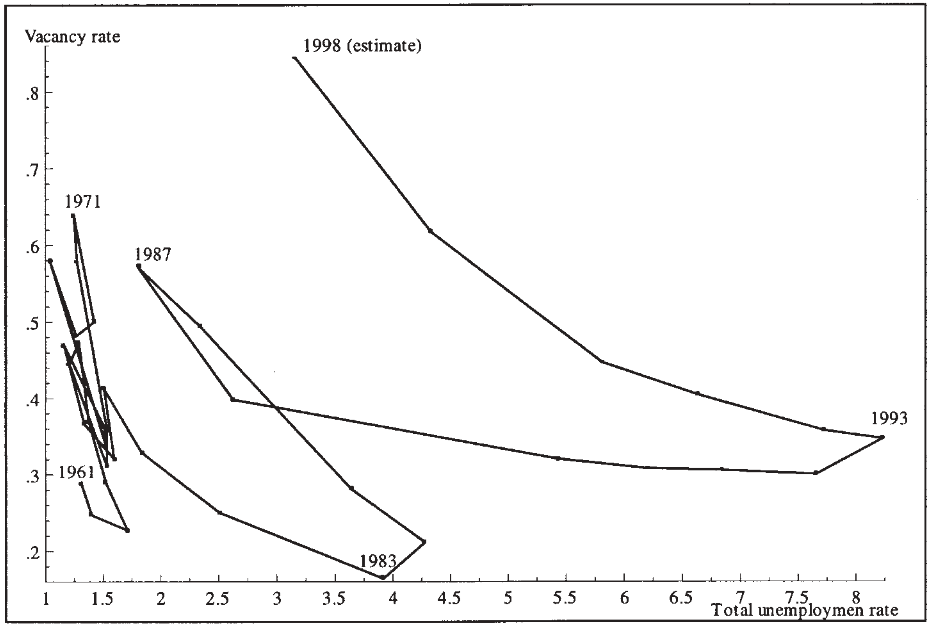

Note: Unemployment rates include open unemployment as well as participants in labour market programs. Sources: The Directorate of Labour (the number of unemployed and the number of vacancies) and Statistics Norway (the size of the labour force).

replacement ratios, high labour taxes, strong trade unions and a very compressed wage distribution. Nevertheless, the Norwegian unemployment experience is somewhat exceptional. From a bottom level of around two per cent in 1987, the unemployment rate rose steadily to eight per cent in 1993, and then fell back again to around three per cent in 1998. In light of the experiences provided by other European countries, the magnitude of the fall in unemployment came as a bit of a surprise. The achievement was partly attributed to the sheer force of the recovery, and partly to the existence of a wide-ranging incomes policy co-operation that kept wage growth in check, at least until 1997. However, many economists considered the low unemployment level in 1998 as a symptom of 'overheat' and therefore unsustainable in the long run. One indication that something had changed was that the number of vacant jobs was much higher in 1998 than what was previously associated with similar levels of unemployment, as illustrated by Figure 1. Moreover, in 1998, the previously successful attempts to control nominal wage growth failed.

In any case, the Norwegian experience offers a welcome opportunity to identify trends in the *composition* of unemployment rates, without having to worry about the 'noise' associated with pure level-effects. We have at our disposal an almost complete cycle, which in terms of aggregate unemployment ended were it started. In this paper, we investigate whether or not this restoration of the 'aggregate' also applied to the underlying micro 'anatomy' of unemployment.

The next section gives a brief description of the data. Section 3 takes a look at changes in the composition of unemployment, in terms of incidence and duration. Section 4 investigates how the overall distribution of unemployment has developed. Section 5 offers a closer look at two different cohorts of unemployed. Section 6 concludes.

2 The data

Our main data source is a large database constructed on the basis of various administrative registers in Norway. The heart of the database is a complete record of all registered spells of unemployment in Norway from January 1989 to March 1998. This information is matched to other registers providing information about age, gender and education. Unemployment spells are divided into open unemployment and participation in various labour market programs. In this paper, we focus on total individual unemployment exposure, regardless of its kind. The length of a spell is determined from a 'starting date' set by the labour office, which is reasonably accurate, and a 'stopping date' set by Statistics Norway (or ourselves) on the basis of the recorded status by the end of each month. Consequently, we only measure unemployment duration in terms of months. Moreover, there are a large number of one-month dropouts from the unemployment register, which we believe are not really associated with the termination of one spell and a start of a new one. In order to be counted as unemployed in Norway, each unemployed person must every fortnight submit a signed card to the labour office, in which availability and willingness to work are confirmed. If a card for some reason is a few days delayed, the unemployment spell is terminated; and when the card finally shows up, it may look like a new unemployment spell in our database. This may happen even if the regular payment of unemployment benefit is completely unaffected (a larger delay is required before these payments are stopped). For that reason, we close one-month gaps by assuming that these secluded short periods outside the unemployment pool are really parts of longer continuous unemployment spells.

In section 4 of this paper, we also use results from five consecutive Level of Living Sample Surveys, conducted by Statistics Norway. These surveys (which were conducted during January-April in 1980, 1983, 1987, 1991 and 1995) contain interview based (retrospective) information about the respondents total unemployment exposure during the previous year, measured in weeks.

3 The incidence/duration composition of unemployment

In order to obtain a crude decomposition of unemployment rates in terms of incidence and duration, we calculate (on the basis of register data) for each month the inflow rate, i_t, (the number of entrants into the unemployment pool divided by the estimated size of the labour force[1]) and the *outflow rate,* o_t, (the number of exits out of the unemployment pool, divided by the number

1. Our data do not contain consistent information about employment relationships. In order to calculate the size of the labour force, we use the employment figures from the Labour Force Sample Surveys (LFSS) together with our own unemployment figures.

Figure 2.
Aggregate stocks and flows. Unemployment in Norway 1989–1998.

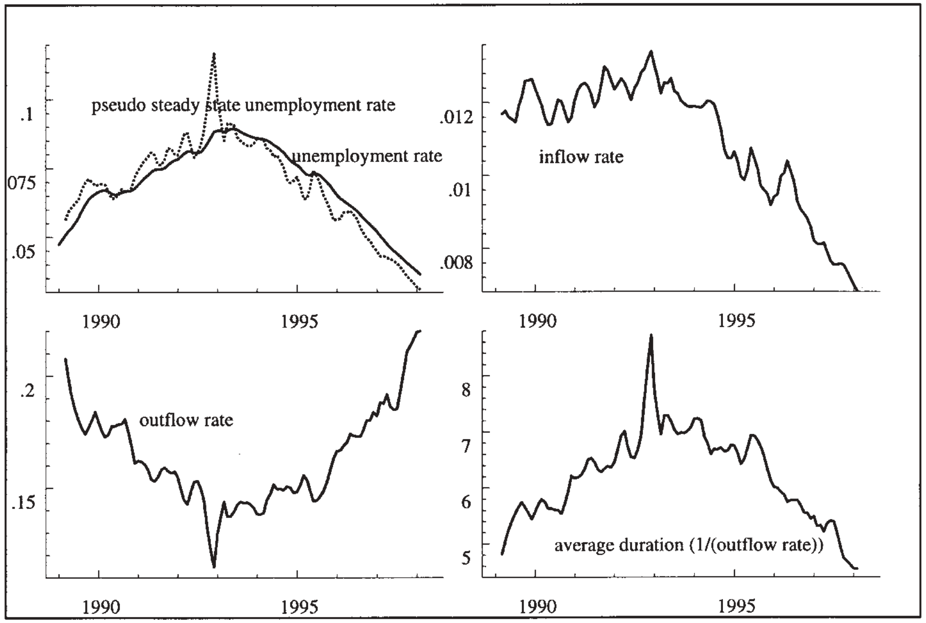

Note: Monthly data. The series are seasonally adjusted and smoothed. Average duration is measured in months. The pseudo steady state level of unemployment is the unemployment level consistent with the currently prevailing (but smoothed) rates of inflow and outflow.

of unemployed in the previous period). We focus on these flows into and out of unemployment, irrespective of where people go to or from. It is likely that the composition of the flows (to/from employment/out of the labour force) changes somewhat over the cycle, although evidence reported by Burda and Wyplosz (1994) for six different countries indicates that the various flow components are surprisingly stable.

There is a marked seasonal component in inflow – and outflow rates (accounting for as much as +/– 20 per cent of the series' level in some months). The inflow rate is typically high during the summer (July-September), as school leavers enter the labour market. It is also relatively high during the coldest part of the winter (November-January), as a number of outdoors activities close down. For similar reasons, the outflow rate is high during the summer (April-September) and low during the winter (October-March). In order to focus on the trend- and cycle components, we remove these seasonal components, and base the discussion in this section on seasonally adjusted and smoothed series[2]. Figure 2 gives

2. Seasonal adjustment was done with the ratio-to-moving-averages method (Makridakis et al, 1983, pp 137-141).

a picture of how aggregate rates of unemployment, inflow and outflow have developed over the past 10 years[3]. We also plot the rate of unemployment towards which the economy would have converged, had the current (smoothed) levels of inflow and outflow continued to prevail (the 'pseudo steady state'). The first thing to note is that the Europe-style ratchet effects in aggregate unemployment rates do not seem to play a major role in Norway. The upwards – and downwards sloping segments of the unemployment level curve are almost mirror images. But does this apparent symmetry also apply to the underlying aggregate flows? The answer is not very far from yes. Figure 2 may convey the impression that the rise in unemployment during the early 1990's was a pure duration phenomenon. But the sharp fall in outflow rates was in fact preceded by a fall in inflow rates during 1988[4]. This is illustrated by the fact that at the start of 1989, the underlying flows already anticipated much of the subsequent increase in the level of unemployment (as indicated by the pseudo steady state). Hence, the rise in unemployment, as well as the subsequent decline, was associated with parallel, although slightly differently timed, changes in incidence and duration. For example, the decline in the pseudo steady state rate of unemployment from its peak level of approximately 10 per cent in late 1992 to 3.5 per cent five years later, was associated with a 43 per cent reduction of the inflow rate (from 13.6 to 7,4 per thousand) and a 41 per cent reduction in average duration (from 8 to 4.7 months). However, in order to identify pure changes in the incidence-duration composition over time, a comparison of inflow – and outflow rates across periods with similar unemployment rates is probably more adequate. And such comparisons do seem to tell a slightly different story. For example, the pseudo steady state rate of unemployment in the beginning of 1996 was almost exactly equal to that prevailing seven years earlier. But the composition of the underlying flows was different; average duration had risen with 25 per cent, while the inflow rate had fallen with 19 per cent.

Now, if a cyclical or structural process of sorting was going on in the labour market – e.g. through discouragement or changes in relative labour demand – this should probably be reflected in the relative performance of various subgroups of the population. Figure 3 reveals that unemployment rates for different groups, classified according to age, gender and education[5], have not at all diverged. On the contrary, there has been some convergence in unemployment rates for men and women,

3. Unsurprisingly, the inflow rate is countercyclical while the outflow rate is procyclical. However, in line with the findings for France, Germany, Spain, UK, USA and Japan reported in Burda and Wyplosz (1994), the gross flows are *both* strongly countercyclical. In fact there is a positive correlation between the gross flows into and out of unemployment of 0.59 (0.76 for the trend-cycle-component).
4. We do not have access to micro data for 1988. However, the stock-numbers published by the Directorate of Labour (number of unemployed less than four weeks) indicate that the inflow rate increased with as much as 80–90 per cent from January 1988 to January 1989.
5. The age- and gender specific rates are calculated on the basis of quarterly total employment estimates, as reported by the LFSS. As our own unemployment files do not provide consistent information about educational attainment for the whole 10-year period, we have not been able to calculate skill-specific rates in the same way. However, we do have reliable information about education for the whole population in 1993. We limit our calculations to workers that can be assumed to have completed their education before 1989 and that are of a typical working age (35–55 in each year). The reported education-specific rates are calculated on the basis of the population, rather than the labour force only.

Figure 3.
Unemployment rates for various subgroups 1989–1997.

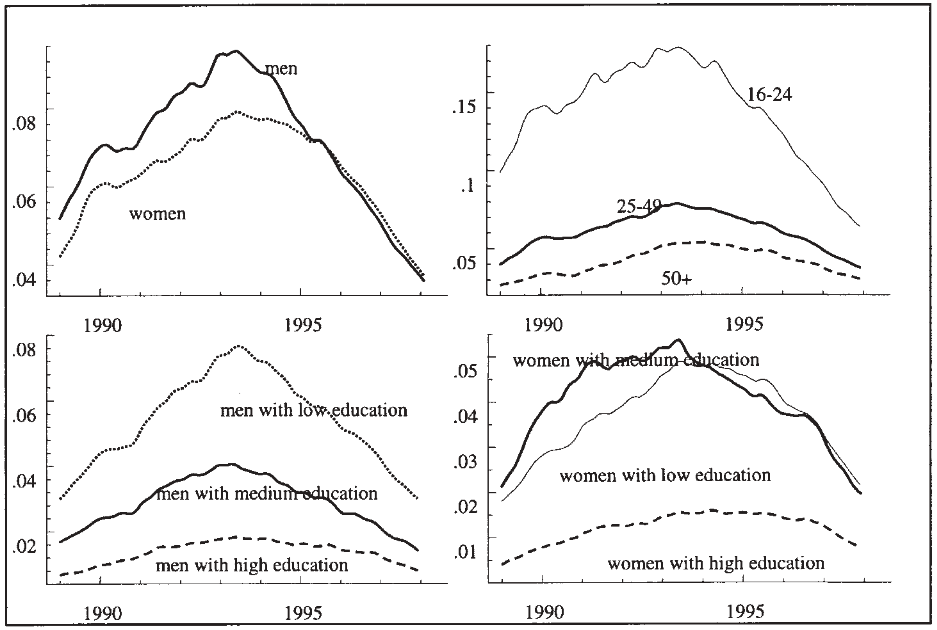

Note: The series are seasonally adjusted and smoothed. The two upper panels display unemployment labour force ratios according to gender and age. The two lower panels display unemployment population ratios for men and women 35–55 years, according to educational attainment. Low education is 10 years or less (primary school), medium education is 11–12 years (completed secondary education), and high education is 13 years or more (college/university).

for young and old, and for women with low and medium education. Figure 4 highlights *relative* changes over time for men and women and for different age groups. The relative positions of young (16–24) and old (above 50) workers have changed markedly. While the relative unemployment rate for the young declined, it rose quite sharply for older workers. In both cases, the change was a pure *duration*-phenomenon. At the start of the period (in 1989), the three age groups (young, middle-aged, old) had quite similar outflow rates (i.e. the relative rates were close to unity). But since then, the relative outflow rate for the young has risen steadily to a level that is approximately 50 per cent above the middle-aged; while the outflow rate for the old has dropped to a level that is 40 per cent below this group. At the same time, their relative inflow rates have remained almost unchanged. Figure 4 also reveals that the relative increase in female unemployment was a pure incidence-phenomenon. The convergence of male – and female inflow rates were probably driven by two forces. First, men typically work in sectors that are sensitive to business cycles;

Figure 4.
Relative rates of inflow, outflow and unemployment according to age and gender, 1989–1997.

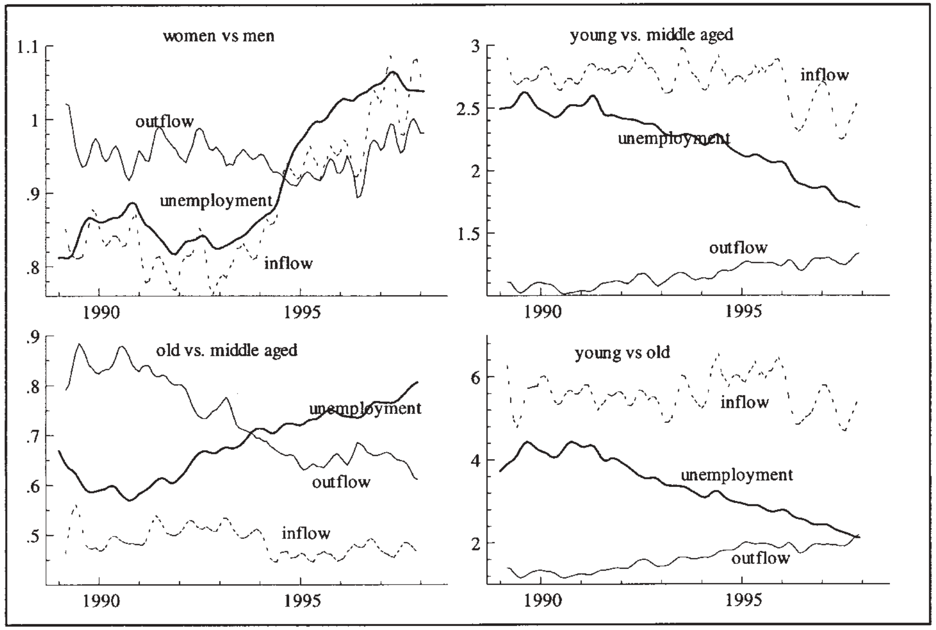

Note: The 'group A vs. group B'-plots display inflow –outflow –and unemployment rates for group A divided by the corresponding rates for group B.

hence their inflow rate was strongly reduced during the economic recovery after 1993. And second, many women seized the opportunity offered by the same recovery to enter the labour market, some of them through a spell of unemployment.

Figure 5 displays the relative performance of the different skill groups. The skill-bias-labour-demand hypothesis suggests that low-skilled workers have confronted an increasing risk of getting stuck in unemployment. Our data do not provide strong support for this hypothesis. For men, we find some (weak) support for the proposition that the relative unemployment rate for low–skilled workers have increased, but this appears to be caused by higher incidence, rather than longer durations. For women, it seems that the relative unemployment rate for the low skilled has increased because of both higher inflow and lower outflow. But these results are likely to reflect changes in the composition of the labour force, rather than changes in the structure of unemployment (women with low education entered the labour market). For both men and women, it appears that the most high skilled workers tend to be relatively more unemployed than before, primarily because of higher incidence. This development is probably related to the large increase in the supply of high skilled workers that resulted from the extraordinary expansion of

Figure 5.
Relative rates of inflow, outflow and unemployment according to educational attainment, men and women (35–55 years), 1989–1997.

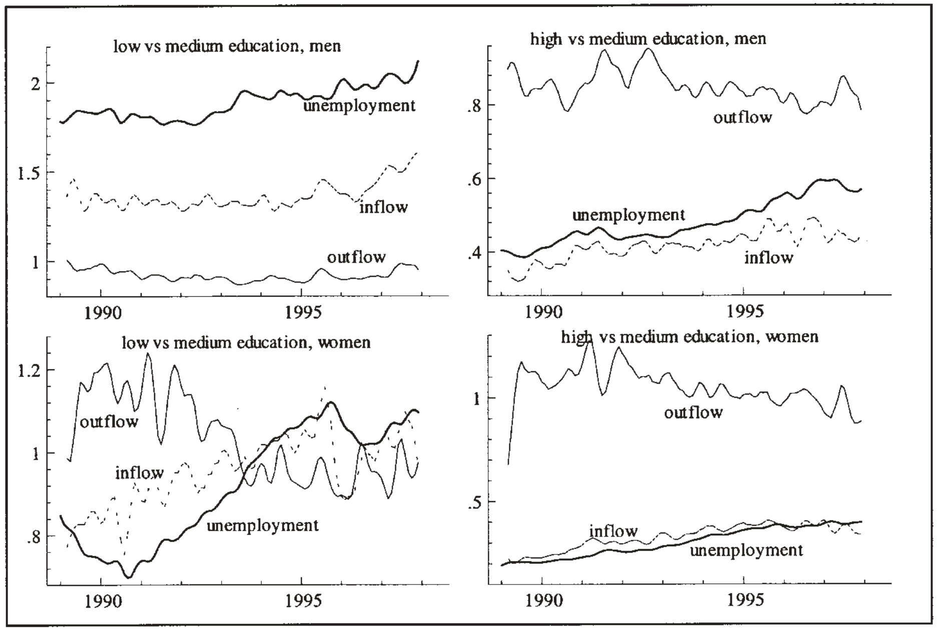

Note: The 'group A vs. group B'-plots display inflow – outflow – and unemployment population ratios for group A divided by the corresponding ratios for group B. The ratios are calculated for men and women aged 35–55 years in each year.

enrolment into universities and colleges during the last recession[6].

Structural changes responsible for increased inequality need not be associated with easily observed individual characteristics. For example, the workers that are most adversely affected by changes in relative labour demand may not be those with low education in particular, but rather workers with low individual productivity relative to others with the same education. A distinctive feature of the European unemployment problem is the very high proportion of long-term unemployed. The high, and to some extent increasing, level of long-term unemployment is sometimes interpreted as indicating that an increasing number of long-term unemployed has become more or less unemployable, and that this phenomenon is one of the driving forces behind the high overall unemployment rates in Europe. Figure 6 displays the outflow rates at four different unemployment dura-

6. The student population in Norway (universities and colleges) increased with 45 per cent from 1988 to 1992. The resultant new academics are only to a small extent represented in our figures, as we have restricted attention to those between 35 and 55 years.

Figure 6.
Outflow rates at different unemployment durations 1989–1998.

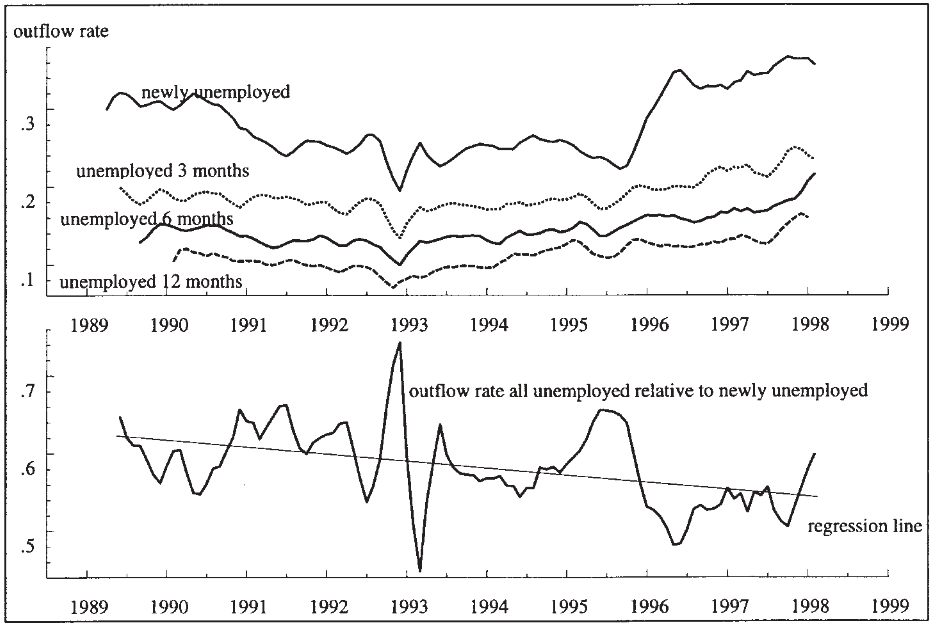

Note: The numbers are seasonally adjusted and smoothed. Newly unemployed are defined as persons with up to one month unemployment.

tions in Norway. It is clear that the outflow rate is unambiguously higher the shorter is unemployment duration. On average, the outflow rate after half a year (one year) of unemployment is 47 (40) per cent of that for newly unemployed. However, relative outflow rates do not seem to have changed very much over time. The rise in unemployment during 1989–1993, as well as the subsequent recovery, was associated with more or less parallel shifts in the outflow rates at all durations of unemployment. The outflow rate for newly unemployed is somewhat more cyclical than the outflow rate at other durations. As indicated in the lower panel of Figure 6, there is also a weak trend in terms of short-term unemployed doing slightly better than before relative to all unemployed[7]. But there is no dramatic (relative) fall in the outflow rate for persons with for example one year of unemployment.

What does the composition of unemployment in Norway look like compared to other countries? Unfortunately, the evidence based on direct observation of flows is sparse. OECD (1990, p. 12) report monthly inflow- and outflow rates for 18 different countries for three different years (1979, 1983, 1988), estimated on the basis of stock figures[8]. For Norway, the estimated 1988-rates of inflow and outflow are 7.9 per thousand and 30.3 per cent respectively. Given that the rate of unemployment was extremely low in 1988 (2.6 per cent), these numbers are not far out of line with our own results. According to the OECD estimates, most European countries have outflow rates that are much below the Norwegian standard. The typical pattern is that European labour markets are characterised by small flows in general, while the North-American labour markets are characterised by large flows. This pattern is confirmed by Current Population Survey data for France (1990) and the United States (1989), reported by Cohen et al (1997, pp 272-273). They find an aggregate monthly rate of inflow[9] of approximately 6–7 per thousand in France and as much as 3 per cent in the United States. The outflow rate is around 5 per cent in France and 25 per cent in the United States. Cohen et al (1997) also report education-specific stocks and flows. In both France and the United States, the low-skill unemployment rate is much higher than the high-skill unemployment rate. But, while this is a pure incidence-phenomenon in the United States (durations tend to be shorter for the low skilled), it is both an incidence – and a duration phenomenon in France. At the aggregate level, Norway seems to have managed to combine French style inflow rates with US style outflow rates. In terms of the skill-structure, Norway is more similar to the United States than to France; higher unemployment among the low skilled is primarily caused by higher incidence.

4 The distribution of unemployment

The incidence-duration decomposition may be misleading if, as is often thought to be the case, the inflow rate is inflated by a large number of re-entrants into the unemployment pool. And indeed, even though we have closed one-month-gaps in the register (see section 2), there are a sizeable number of persons that move into and out of the unemployment register intermittently. These persons contribute to a large number of inflows as well as outflows, even though many of them may actually be unemployed all the time. It is possible that the re-entrants conceal changes in the true composition of incidence and duration; or more generally, in the distribution of unemployment.

In this section, we take a different approach, and consider how *total unemployment exposure* in different calendar time periods was distributed. Intuitively, the degree of inequality in the unemployment distribution depends on three factors; the fraction of people exposed to at least some unemployment, the total volume of unemployment for those that are exposed, and the distribution of that total volume. But unfortunately,

7. Jackman and Nickell (1991) show that the outflow rate for all unemployed relative to newly unem-ployed may be used to detect duration dependence in individual exit probabilities. In the absence of duration dependence, this ratio is unaffected by the level of unemployment. The argument builds on comparisons of steady states; hence its empirical relevance with respect to the pattern depicted in the lower panel of Figure 6 is questionable.

8. The inflow rate is calculated as the number of persons unemployed for less than a month, divided by the working age population (15–64) less the unemployed. The outflow rate is calculated as the difference between the average monthly level of inflows and the average monthly change in unemployment over the whole year, divided by the average number of unemployed.

9. In their paper, the inflow rate is related to the flow from employment to non-employment, while the outflow rate is related to the flow from unemployment to employment.

standard relative inequality-measures such as Lorenz curves and Gini-coefficients fail to pick up all these elements. To illustrate, let p_{0t} be the fraction of the labour force in period t with zero unemployment. Let μ_t be the mean and G_t be the Gini-coefficient associated with the distribution of unemployment within the whole labour force (including those with zero unemployment), and let $\overline{\mu}_t$ and $\overline{G}_t$ be the same measures calculated for those with non-zero unemployment only. We then have that

$$G_t = \overline{G}_t + p_{0t}(1-\overline{G}_t),$$

hence the 'total Gini' captures two of the elements referred to above; the inequality among the unemployed (the 'partial Gini') and the fraction of persons exposed to unemployment. But it does not capture the degree of inequality *between these two groups*. If for example the rate of unemployment doubles such that the entire increase in unemployment exposure is allocated proportionally to those that are already exposed to some unemployment, the two Gini measures remain completely unchanged. The reason of course, is that these measures are invariant with respect to changes of scale. This is typically viewed as a desirable property for inequality-measures aimed at evaluating income distributions. But it may not be equally desirable for measures aimed at evaluating unemployment distributions, for which the difference between those that never become unemployed (the zeros) and those that do (the positive quantities), plays a key role. An alternative to the requirement of invariance towards equal proportional changes is the requirement of invariance towards equal absolute changes. One measure that remains unchanged if all workers in the labour force are exposed to the same additional amount of unemployment is the 'absolute Gini', $G_t^a = \mu_t G_t$. The Gini-coefficients have very simple interpretations. Let Δ_t be the absolute value of the difference in total unemployment exposure between two persons randomly selected from the labour force, and let $\overline{\Delta}_t$ be the corresponding difference for two persons randomly selected from the group of persons with at least some unemployment. We then have that $E\Delta_t = 2G_t^a = 2\mu G_t$ and $E\overline{\Delta}_t = 2\overline{G}_t^a = 2\mu\overline{G}_t$. Hence, the absolute Gini-coefficient (multiplied by two) measures the expected absolute difference in unemployment between two randomly selected individuals, while the relative Gini-coefficient measures the same expected difference relative to the mean.

In reality, changes in the aggregate level of unemployment cannot be distributed equally according to any of the invariance requirements. The reason is that unemployment exposure for a single person during a given period of time is bounded between zero and the length of the period under consideration. Therefore, the maximum values of the Gini-coefficients change from period to period (and are never equal to one), depending on the total level of unemployment. One solution to this problem is to evaluate the Gini-coefficients relative to their period-specific maximum (given the total level of unemployment). These adjusted Gini-indexes are the same for the relative and absolute measures. Let $(p_{0t}^{\max}|\mu_t)$ be the fraction of persons that would have had zero unemployment in the hypothetical case in which the given total level of unemployment was distributed on as few persons as possible. The adjusted total Gini, G_t^* is given by:

$$G_t^* = \frac{\overline{G}_t + p_{0t}(1-\overline{G}_t)}{p_{0t}^{\max}|\mu_t} = \frac{\overline{G}_t + p_{0t}(1-\overline{G}_t)}{1-\frac{1}{k}\mu_t},$$

where k is the maximum level of individual unemployment exposure, i.e. the length of the period under consideration. A similar measure may be calculated for the partial

Gini, for which maximum inequality is obtained when all the unemployed persons have either the minimum or the maximum degree of unemployment.

Another inequality-measure that is invariant towards equal absolute changes for all is the empirical variance, V_t. The empirical variance has the additional advantage that it is linearly decomposable. Assume for example that we divide the population into J groups, $j=1,2,...J$, and that we want to know how much of total variance is accounted for by the difference between these groups. Denote the within-group variances by V_{jt}, and the between-group variance by V_{Bt}. The latter is calculated as the variance that would have resulted if all persons in each group were characterised by the group-mean. We then have that

$$V_t = \sum_j f_{jt} V_{jt} + V_{Bt},$$

where f_{jt} is the fraction of the population belonging to group j in period t. The fraction of total variance accounted for by differences between the J groups is accordingly

$$1 - \sum f_{jt} V_{jt} / V_t.$$

Rather than relying on one single inequality-measure, we report below a number of alternative measures, such as various Gini-coefficients, variances, and other key properties of the distribution that illuminate the degree of inequality from different angels. In order to calculate these measures, two potentially consequential decisions must be made. First, an appropriate time span must be selected and, secondly, a way to estimate the relevant number of persons without any unemployment exposure during that period must be found. We approach these difficulties in two alternative ways. In the first approach, we select calendar years as the unit of measurement, and calculate the number of zero-unemployment-observations from the labour force estimates reported in the LFSS. Obviously, total unemployment exposure for a single individual during a year cannot exceed 12 months, hence total unemployment exposure for the most unemployed does not (almost by definition) change very much from year to year. In the second approach, we select three-year-periods as the unit of measurements, and focus on a subgroup of the population for which the labour force participation rate is very high, namely prime aged men (35–55 years). By assuming that all these persons are exposed to the risk of unemployment throughout the period, we can easily calculate the number of zero-observations for each of the three three-year-periods.

Table 1 and Table 2 report the various inequality-measures on a yearly and a three-yearly basis, respectively. The message conveyed by Table 1 is that the overall degree of inequality in the unemployment distribution has increased, but not very much. In 1989, 15 per cent of the labour force were unemployed for an average number of 4.56 months each. In 1997, only 12 per cent of the labour force were exposed to unemployment, but the average number of months had risen to 4.68. Moreover, the fraction of unemployed being unemployed throughout the whole year rose sharply (with 35 per cent), despite the fall in aggregate unemployment from 5.66 to 4.50 per cent. The table reveals that the crude Gini-coefficients are highly correlated to the *level* of unemployment. The relative Gini-coefficients are strongly pro-cyclical, while the absolute Gini-coefficients are counter-cyclical, indicating that when e.g. unemployment rises, the additional volume of unemployment is distributed more equally than the alternative of equal proportional increases,

Table 1.
The Yearly Distribution of the Total Number of Unemployment Months 1989–1997. Whole Labour Force.

	1989	1990	1991	1992	1993	1994	1995	1996	1997
a)									
Total number of persons in the labour force (1000)	2165	2149	2135	2139	2143	2161	2197	2239	2285
Total number of unemployed during the whole year (1000)	322	347	369	390	409	394	357	321	264
Total number of unemployment months (1000)	1469	1715	1893	2104	2266	2140	1907	1605	1233
Per cent of labour force unemployed simultaneously (average monthly unemployment rate)	5.66	6.65	7.39	8.20	8.82	8.26	7.23	5.98	4.50
Per cent of labour force having some unemployment during the year	14.85	16.13	17.22	18.25	19.04	18.10	16.10	14.32	11.54
Per cent of unemployed being unemployed at least 6 months during the year	32.66	37.30	39.37	41.07	43.27	42.18	41.37	37.81	34.12
Per cent of total unemployment accounted for by those with at least 6 months	62.75	66.59	69.09	70.78	72.53	71.47	70.99	68.77	65.27
Per cent of unemployed being unemployed the whole year	6.89	9.19	10.94	13.72	14.72	13.83	13.41	11.59	9.33
Per cent of the unemployed having 50 per cent of all unemployment	23.39	24.44	24.46	24.64	25.05	24.99	24.70	23.54	22.89
Average number of unemployment months in labour force	0.68	0.80	0.89	0.98	1.06	0.99	0.87	0.72	0.54
Average number of unemployment months for those having some unemployment	4.57	4.95	5.14	5.39	5.55	5.46	5.38	5.01	4.68
Average number of unemployment months for the 25 per cent most unemployed	9.59	10.05	10.46	10.89	11.08	10.93	10.85	10.48	9.94
b)									
Partial Gini (Gini for the unemployed)	0.409	0.393	0.392	0.390	0.383	0.384	0.387	0.408	0.415
Absolute partial Gini	1.867	1.943	2.018	2.099	2.126	2.100	2.083	2.041	1.942
Total Gini (Gini for the whole labour force)	0.912	0.902	0.895	0.889	0.883	0.889	0.901	0.915	0.933
Absolute total Gini	0.619	0.720	0.794	0.874	0.934	0.880	0.782	0.656	0.503
Adjusted partial Gini	0.775	0.768	0.782	0.796	0.798	0.792	0.792	0.801	0.793
Adjusted total Gini	0.967	0.966	0.967	0.968	0.968	0.969	0.972	0.973	0.976
c)									
Total Variance (in whole labour force)	4.355	5.277	6.047	6.889	7.480	6.993	6.179	5.014	3.676
Total Variance relative to Maximum Variance (given total unemployment exposure)	0.567	0.590	0.614	0.636	0.646	0.641	0.639	0.620	0.594
Per cent of variance accounted for by the difference between positive and zero observations	85.15	83.87	82.70	81.73	80.91	81.77	83.77	85.68	88.46
Per cent of variance accounted for by the difference between men and women	0.08	0.05	0.09	0.11	0.11	0.04	0.00	0.00	0.00
Per cent of variance accounted for by the difference between age groups	3.11	5.06	3.52	3.20	2.85	2.35	1.69	1.11	0.62
Per cent of variance accounted for by the difference between educational groups	0.90	1.02	1.08	1.15	1.18	1.03	0.88	0.72	0.55

Note: The adjusted Gini coefficients are the crude coefficients divided by their value under maximum possible inequality, given that no one can be unemployed for more than 12 months.

Table 2.
The Three-year Distribution of the Total Number of Unemployment Months 1989-1997. Prime Aged Males (35–55 years).

	1989–1991	1992–1994	1995–1997
a)			
Total population (1000)	517	557	592
Total number of unemployed during the three year period (1000)	75	102	78
Total number of unemployment months (1000)	672	1060	748
Per cent of labour force unemployed simultaneously (average monthly unemployment rate)	3.61	5.28	3.51
Per cent of labour force having some unemployment during the three year period	14.47	18.31	13.18
Per cent of unemployed being unemployed at least half the period	16.82	21.95	18.76
Per cent of unemployed being unemployed the whole period	0.62	2.60	1.58
Per cent of the unemployed having 50 per cent of all unemployment	19.40	21.57	18.42
Average number of unemployment months in labour force	1.30	1.90	1.26
Average number of unemployment months for those having some unemployment	8.98	10.39	9.58
Average number of unemployment months for the 25 per cent most unemployed	19,93	23,38	20,80
b)			
Partial Gini (Gini for the unemployed)	0.492	0.502	0.500
Absolute partial Gini	4.420	5.220	4.787
Total Gini (Gini for the whole labour force)	0.927	0.909	0.934
Absolute total Gini	1.204	1.729	1.179
Adjusted partial Gini	0.717	0.760	0.739
Adjusted total Gini	0.961	0.960	0.968

Note: The analysis encompasses, for each three-year-period, men that are at least 35 in the first year and not more than 55 in the last.

Figure 7.
Adjusted Gini-coefficients for the yearly distribution of total unemployment exposure 1989–1997.

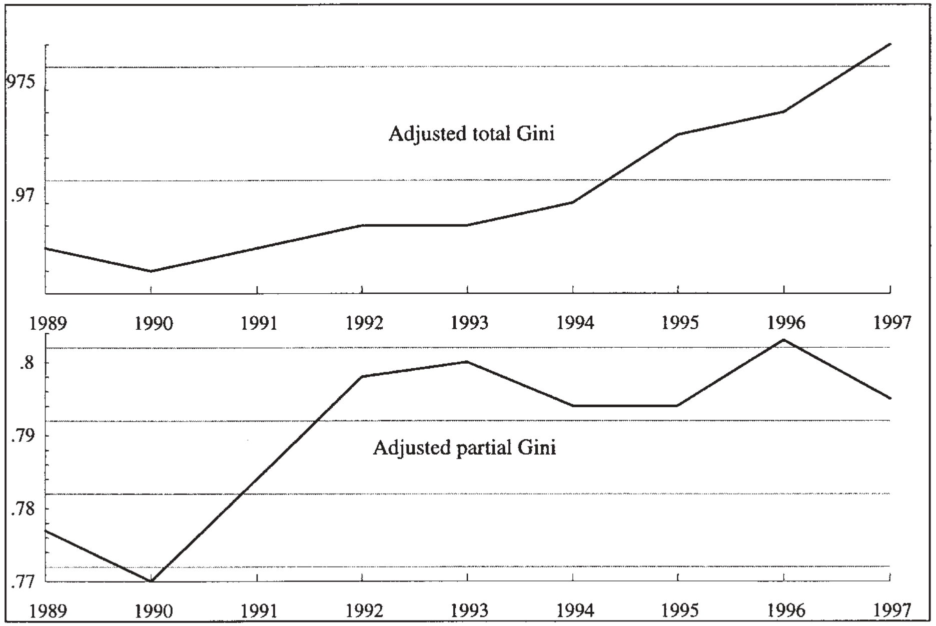

Note: The adjusted Gini coefficients are the ordinary Gini coefficients divided by their hypothetical maximum in each year. The total Ginis are uncoditional, while the partial Ginis are conditional on having at least some unemployment.

but less equally than the alternative of equal absolute increases. A desired property of inequality measures used to compare distributions of unemployment at different points in time is that they are not too closely related to the total level of unemployment. The only measures in the table that stand up to this test are the adjusted Gini-coefficients. These inequality measures do not exhibit any clear cyclical pattern, while they indicate a small, but steady rise in inequality over time[10]. This development is highlighted in Figure 7.

10. To evaluate the degree of cyclical sensitivity, we may consider how an additional month of unemployment affects the yearly Gini coefficients, depending on how it is allocated. Let the additional month be allocated to a person who has j months of unemployment from before, and let p_j be the fraction of workers with j months or less. It can be shown that the total Gini increases iff $2p_j > 1+G$, while the adjusted total Gini increases iff $2p_j > 1+G\,((12-2\mu)/(12-\mu))$. In 1989, for example, an additional month of unemployment would have increased the total Gini only for $p_j \geq 0.956$ or $j \geq 5$. And since those with 12 months of unemployment cannot possibly have any more, the Gini coefficient increases only if the additional month is allocated to the small fraction of the labour force (approximately 3 per cent) having between 5 and 11 months. The adjusted Gini would in this case have increased for $p_j \geq 0.928$ or $j \geq 3$.

Figure 8.
Total unemployment exposure in 1989 and 1996. Relative frequencies.

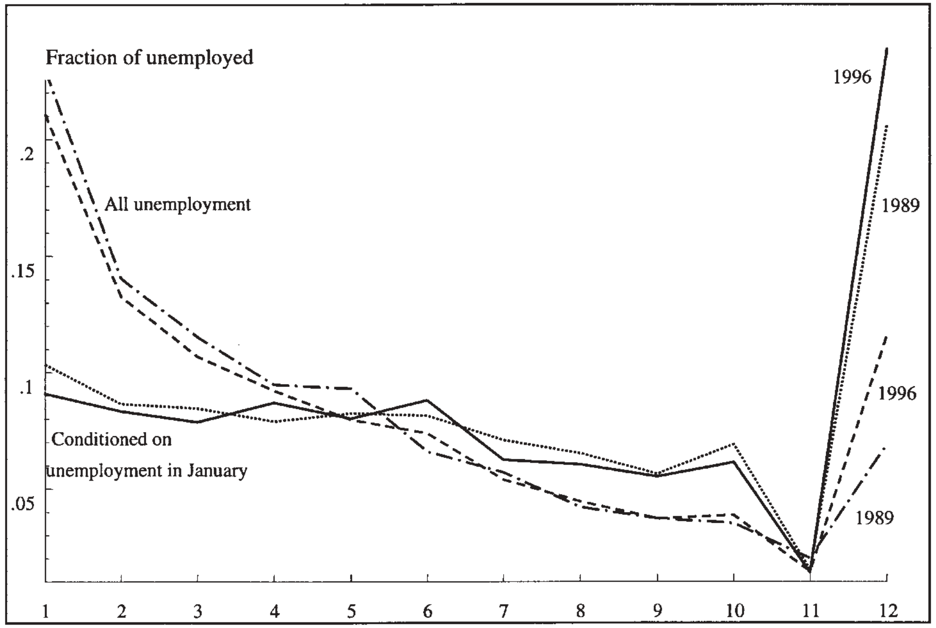

Note: The number of persons exposed to unemployment in 1989 was 321638, of which 107471 were unemployed in January. The corresponding numbers for 1996 were 320626 and 155807.

The difficulty associated with separating level-effects from structural trends suggests that it may be preferable to compare years with similar levels of aggregate unemployment. The best candidates for such comparisons are 1989 and 1996. But although these years were similar in terms of the unemployment level, 1989 was characterised by rising unemployment, while 1996 was characterised by falling unemployment; hence it is difficult to sort out phase-of-the-cycle effects from underlying structural trends.

A comparison of 1989 and 1996 reveals that inequality has increased according to almost all the inequality measures reported in Table 1. The complete unemployment exposure distributions for these two years are depicted in Figure 8. They indicate that the increasing inequality is almost entirely due to a rise in 'permanent' unemployment. While 6.9 per cent of the unemployed (1 per cent of the labour force) were unemployed throughout the year in 1989, this happened to 11.6 per cent of the unemployed (1.6 per cent of the labour force) in 1996. To some extent, this reflects that the stock of unemployment was larger in January 1996 than in January 1989 (7 versus 5 per cent), hence more people were exposed to the risk of being permanently unemployed in the latter year. Therefore, we also plot the two unemployment distributions, conditioned on being unemployed in the first month of each year. The message is similar although less conspicuous. While 20.6 per cent of these workers remained unemployed throughout the year in 1989, this happened to 23.8 per cent in 1996, even though the level of unemployment rose during the first and fell during the latter of these years.

The variance decomposition reported in the lower part of Table 1 indicates that the fraction of total variance explained by age has declined steadily. The fractions explained by gender and education have also declined, but these fractions seem to be more volatile with respect to the cycle. In particular, it seems that educational attainment is responsible for much more (absolute) inequality during a recession than during a recovery. The main message is that if inequality in unemployment exposure has widened during the past ten years, this is not driven by forces that are closely related to age, gender or education.

The three-yearly prime-aged-men inequality-measures offered in Table 2 also indicate rising inequality. Again, the most fruitful comparison to make is that between the first and the last period, as these two periods are similar in terms of the average level of unemployment. In particular, it is worth noting that the fraction of 'permanently' unemployed (i.e. unemployed throughout the whole three-year period) almost tripled from the first to the last period, even though average unemployment exposure fell. The complete frequency distributions plotted in Figure 9 reveal that this is in fact the only significant change in the distribution of unemployment exposure from 1989–91 to 1995–97. If we simply remove the extra 800 permanently unemployed persons from the 1995–97 data-file, the two resultant distributions are hardly distinguishable.

Structural changes in the distribution of unemployment are probably sluggish; hence a time span of ten years may be too short to discover substantial changes. In order to assess how the distribution of unemployment has evolved over a longer period of time, we use data collected from the Level of Living Sample Surveys, starting in 1979 (see section 2). Table 3 contains the main results[11]. It may be of particular interest to compare 1979 and 1986, as these two years are very similar in terms of aggregate unemployment, although they are separated by a slump during 1982–1985. The results indicate that nothing of structural importance happened from 1979 to 1986. The various inequality-measures give divergent results[12].

To our knowledge, there is not much international evidence available with respect to the total distribution of unemployment. The seminal work by Clark and Summers, 1979, offers a sort of benchmark, based on the US unemployment distribution in 1974. In that year, an aggregate unemployment rate of 5.5 per cent was generated by 15 per cent of the labour force being unemployed for an average of 3.5 months. This is not very different from Norway in the 1990's. However, Gini-coefficients reported by Butler and McDonald, 1986, indicate that unemployment was in fact more unequally distributed in the United States in the 1970's than it is now in Norway. While the partial (relative) Gini-coefficients for Norway vary between 0.32 and 0.42, they varied between 0.44 and 0.54 in the United States. Jensen and Jensen (1997) have calculated Gini-coefficients for a panel of Danish workers for the years 1986-90. The yearly corrected total Gini-coefficients range from 0.90 to 0.95, indicating slightly less inequality than in Norway. However, all these various results are very sensitive to the precise definition of unemployment, the unit of measurement (using months as the unit of measurement obviously implies that some short spells are lost, while others are overvalued), as well as to the way data are collected. For example, our results indicate less relative inequality than what was found for Norway by Berg and Børing, 1997, building on a different definition of unemployment (unemployment measured in weeks and including part-time unemployment). They reported (partial) Gini-coefficients ranging from 0.41 to 0.47.

Machin and Manning (1998) report the fraction of total unemployment accounted for by those in it for more than half the year for Australia (1985), Denmark (1980), Sweden (1983), United States (1990), United

11. The 1979-data have previously been used to evaluate the degree of inequality in the unemployment distribution by Andersen and Aaberge (1983).
12. Note that our register data indicate that the Level of Living Sample Surveys seriously underestimate the degree of unemployment. For example, in 1994, 10.8 per cent of the respondents declared that they were unemployed at some time during the year, while the register tells us that the true number was at least 18.1 per cent (the definition of unemployment in the survey is broader than the register definition). However, comparisons of the conditional distributions (for the unemployed) do not reveal substantial differences between the registers and the surveys.

Figure 9.
Total unemployment exposure for men 35–55 during 1989–91 and 1995–97. Relative frequencies.

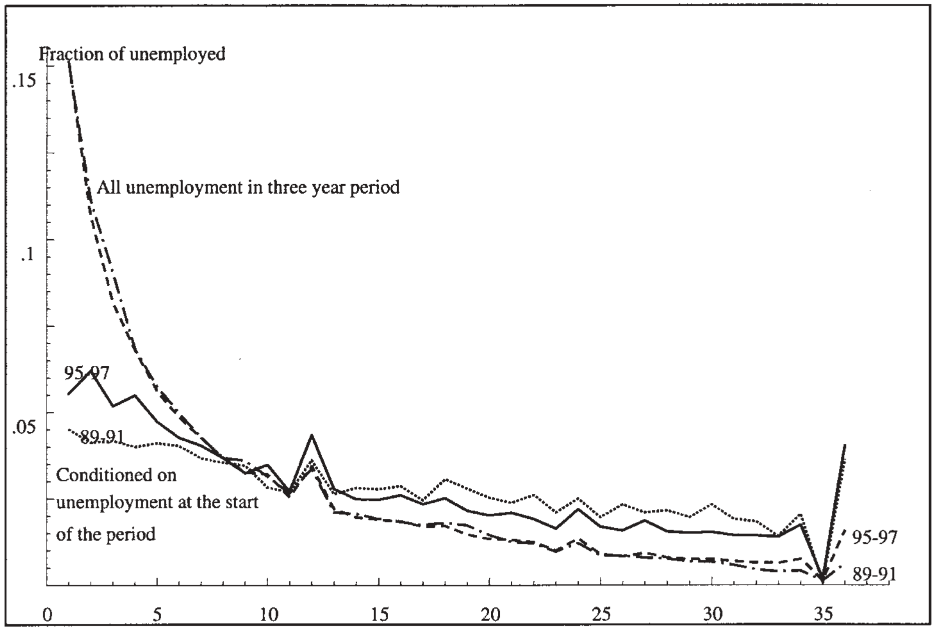

Note: The number of persons (in the relevant group) exposed to unemployment was 74843 during the 1989–91period, out of which 12778 were unemployed at the beginning of the period. The corresponding numbers for the 1995–97 period were 78046 and 30644.

Table 3.
The Yearly Distribution of Total Number of Unemployment Weeks 1979, 1982, 1986, 1990, 1994. Level of Living Survey Data.

	1979	1982	1986	1990	1994
a)					
Total sample	3885	3929	4373	3755	3571
Number of unemployed	281	300	336	410	384
Per cent unemployed simultaneously (unemployment rate according to LFSS)	2.0	2.6	2.0	5.2	5.4
Per cent of labour force having some unemployment during the year	7.23	7.64	7.68	10.92	10.75
Average number of unemployment weeks in whole sample	0.73	0.77	0.71	2.01	1.82
Average number of unemployment weeks for those having some unemployment	10.03	10.11	9.29	18.39	16.96
Per cent of unemployed being unemployed at least 26 weeks during the year	10.32	10.00	8.63	29.76	26.82
Per cent of total unemployment accounted for by those with at least 26 weeks	41.59	35.33	37.28	69.89	64.47
Per cent of unemployed being unemployed the whole year	2.85	0.67	2.98	14.15	9.38
Per cent of the unemployed having 50 per cent of all unemployment	14.23	16.67	13.69	18.29	18.23
Average number of unemployment weeks for the 25 per cent most unemployed	26.63	25.33	25.39	46.39	41.84
b)					
Partial Gini (Gini for the unemployed)	0.556	0.510	0.577	0.526	0.515
Absolute partial Gini	5.571	5.163	5.356	9.670	8.731
Total Gini (Gini for the whole labour force)	0.968	0.963	0.967	0.948	0.948
Absolute total Gini	0.702	0.743	0.690	1.904	1.729
Adjusted partial Gini	0.753	0.693	0.775	0.847	0.799
Adjusted total Gini	0.982	0.977	0.981	0.987	0.983
c)					
Total Variance (in whole labour force)	17.228	15.799	16.389	68.329	56.578
Per cent of variance accounted for by the difference between positive and zero observations	61.021	54.500	62.846	51.971	51.310
Per cent of variance accounted for by the difference between men and women	0.000	0.000	0.000	0.000	0.000
Per cent of variance accounted for by the difference between age groups	0.740	1.785	1.438	1.938	0.855
Per cent of variance accounted for by the difference between educational groups	0.001	0.064	0.251	0.255	0.517

Kingdom (1990) and Germany (1990)[13]. These fractions are lowest in United States (45 per cent), Sweden (53 per cent) and Denmark (63 per cent), and highest in United Kingdom (78 per cent), Germany (76 per cent) and Australia (75 per cent). Our own numbers reported in Table 1 and Table 3 indicate that Norway is in the high-inequality league according to this (conditional) measure. However, the measure's sensitivity with respect to the total level of unemployment, as well as the differences in the way data are collected across countries, suggest that one should be very careful in interpreting these kinds of international comparisons.

5 A tale of two similar unemployment cohorts

At the aggregate level, there is a striking similarity between unemployment statistics during the first halves of 1990 and 1996, although the inequality analysis in the previous section suggests that the composition may have changed. This section takes a closer look at the microanatomy of unemployment in these two periods. We concentrate on two cohorts of unemployed: Those that became unemployed in January 1990 (before the worst part of the recession) and those that became unemployed exactly six years later (after the recession had passed). We first take a look at the crude composition of these two cohorts, in terms of observed characteristics. We then perform unemployment duration analyses for each of the two cohorts, in order to distinguish any changes in the way various individual characteristics affect relative outflow rates or in the degree of duration dependence (or unobserved heterogeneity).

We apply a flexible baseline model with proportional hazards, as recommended by e.g. Meyer (1990) and Narendranathan and Stewart (1993). However, in order to eliminate potential bias arising from changes in the macro-economic conditions and seasonal variability in outflow rates, we estimate individual hazard rates relative to aggregate outflow rates for each month (i.e. we assume that a k per cent increase/decrease in the aggregate outflow rate causes, ceteris paribus, a k per cent increase/decrease in individual hazard rates). More precisely, the hazard models are of the form:

$$\frac{\theta_{ji}(\tau)}{o_{j\tau}} = \exp(\lambda_{j\tau} + x_{j\tau i}'\beta_j), \quad j=1990,1996, \quad i=1,2,\ldots,N_j$$

$$j=1990,1996,\ i=1,2,\ldots,N_j\ ,$$

where j denotes the two cohorts, $o_{j\tau}$ is the aggregate outflow rate (for all cohorts of unemployed) in the calendar-month corresponding to duration-month τ for cohort j, λ_{jp} is the duration-month-specific constant term, $x_{j\tau i}$ is the vector of explanatory variables for individual i (which may change during the spell) and ($\lambda_{j\tau}$, β_j) are the parameters to be estimated. Each cohort is followed for 25 months, after which still ongoing unemployment spells are censored.

Table 4 describes the composition of the two cohorts in terms of some observed variables. We also report some statistics that describe the 'survivors', i.e. those that are still unemployed by the end of the two-year observation period. Figure 10 displays the estimated relative hazard rates for each of the 25 duration-months, normalised on the hazard for the first month (with 95 per cent confidence intervals). Both hazards display a

13. The first three are gathered from OECD (1985), while the latter three are based on the authors own calculations.

Table 4.
The Composition of the Two Cohorts.

	The January 1990-cohort		The January 1996-cohort	
	The whole cohort	Still unemployed after 2 years	The whole cohort	Still unemployed after 2 years
Number of persons	37678	1312	34851	700
Number of unemployment months	211748		179631	
Average duration (months)	5.6		5.2	
Constant covariates				
Gender (per cent)				
Men	61.96	63.30	57.11	49.86
Women	38.04	36.70	42.89	50.14
Age (per cent)				
16–19	12.98	5.58	7.54	2.14
20–24	23.42	21.67	20.59	6.00
25–49	54.17	55.94	60.52	57.43
50–59	6.55	7.11	8.37	16.29
Above 60	2.88	9.70	2.98	18.14
Years of education (per cent)				
9 or less	25.01	34.92	19.71	30.86
10	26.50	31.93	22.34	24.14
11–12	33.92	24.74	44.11	32.43
13–16	5.68	4.53	10.51	10.57
17 or more	0.51	0.65	2.92	2.00
Unknown	8.37	3.23	0.41	0.00
Experience (average)				
Unemployment last year (average number of months)	1.54	2.10	1.22	1.53
Previous work experience (average number of years)	6.50	8.96	7.77	10.16
Country /language background (per cent)				
Not speaking Norwegian	2.44	5.98	7.86	17.43
Non-OECD citizenship	3.53	7.76	5.01	11.57
Previous state (per cent)				
Employment	68.40	65.89	68.76	66.86
Out of the labour force	31.60	34.11	31.24	33.14
Time-varying covariates				
Participation in labour market programs (per cent)				
Average across individuals of average across duration months	6.58		4.27	
Total unemployment months	34.14		28.34	
Participating at some time in the year prior to- or during the spell	44.41		21.63	
Unemployment benefits/program support (per cent)				
Average across individuals of average across duration months	51.87		50.09	
Total unemployment months	83.25		74.76	

Figure 10.
Estimated relative hazard rates during the first 25 months of unemployment spells (with 95 per cent point-wise confidence interval).

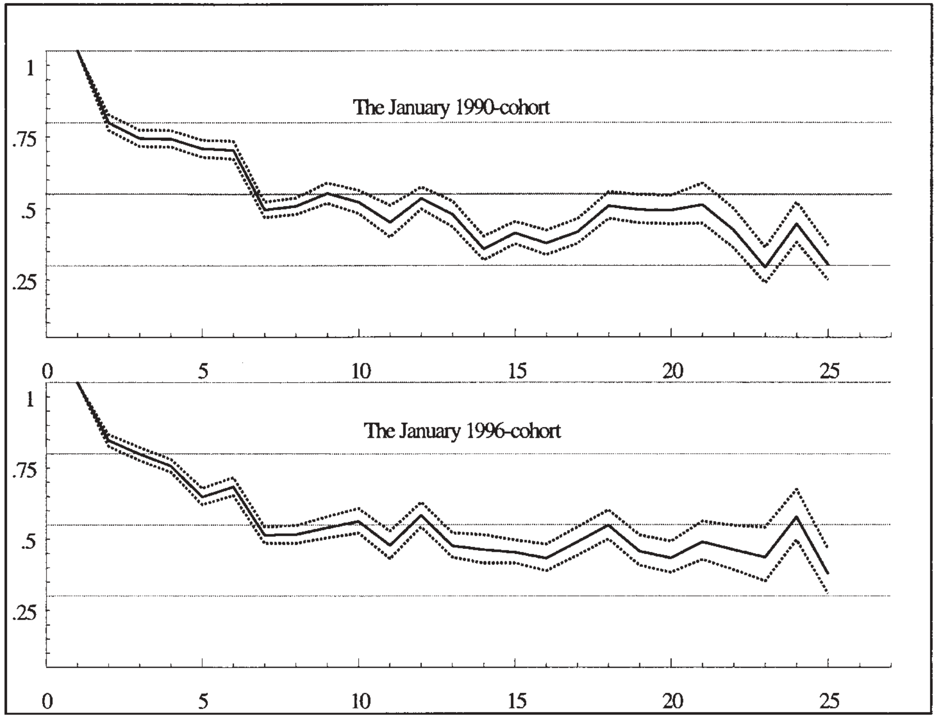

Note: The figures display the estimated hazard rates for each duration-month, relative the hazard for period 1.

pattern of sharply (and almost monotonously) declining exit probabilities. This pattern may result from the existence of unobserved heterogeneity and/or negative duration dependence at the individual level; and whatever its cause, the phenomenon seems to be more manifest in 1990 than in 1996. Table 5 offers some estimation results[14] for the constant and time-varying covariates. The estimates suggest that workers with low education actually improved their relative performance slightly from 1990 to 1996, hence the skill-biased labour demand hypothesis (in terms of educational attainment) does not receive much support from these results either. The hazard rates for women and

14. The estimations also included a large number of covariates for which we do not report estimated co-efficients, among them county of living, profession and previous income (a total number of 79 covariates). Detailed results are available on request.

Table 5.
Hazard Rate Estimates and Relative Risks. Flexible Baseline, Proportional Hazards.

Selected covariates	The January 1990-cohort Coefficient Estimate	Standard Error	Relative Risk	The January 1996-cohort Coefficient Estimate	Standard Error	Relative Risk
Men	-0.08	0.02	0.93	-0.01	0.02	0.99
Women			1.00			1.00
Married	0.20	0.02	1.21	0.10	0.02	1.11
Married and women	-0.23	0.02	0.79	-0.09	0.02	0.92
Number of children	-0.03	0.01		-0.03	0.01	
Age 16–19	0.06	0.02	1.06	0.17	0.02	1.18
20–24	-0.07	0.02	0.94	0.17	0.02	1.19
25–49			1.00			1.00
50–59	0.02	0.02	1.02	-0.20	0.02	0.82
Above 60	-0.32	0.03	0.73	-0.60	0.04	0.55
9 years of education or less	-0.12	0.02	0.88	-0.07	0.02	0.93
10 years	-0.11	0.01	0.90	-0.08	0.01	0.93
11–12 years			1.00			1.00
13–16 years	0.05	0.03	1.06	0.00	0.02	1.00
17 years or more	-0.06	0.08	0.94	0.05	0.03	1.06
Previously employed			1.00			1.00
Previously out of labour force	-0.12	0.01	0.88	-0.12	0.01	0.89
Number of months unemployed previous year	-0.04	0.00		-0.04	0.00	
Program participation	-0.60	0.02	0.55	-0.72	0.02	0.49
Previous program experience	0.31	0.02	1.36	0.22	0.02	1.25
Benefits/program payments	-0.28	0.01	0.76	-0.14	0.01	0.87
Not speaking Norwegian	-0.18	0.04	0.83	-0.23	0.03	0.79
Non-OECD citizenship	-0.21	0.03	0.81	-0.01	0.04	0.99

men also became more equal. In the 1990-cohort, unmarried women had a higher and married women a much lower exit probability than similar men. Even though this pattern also appears in the 1996-cohort, it is clearly less pronounced. The by far most conspicuous difference between the two cohorts is found in the relative performance of different age groups. While the youngest did much better in 1996 than in 1990, the exit rates of unemployed above 50 years deteriorated sharply. The relative hazard rate for those between 50 and 59 fell from 1.02 to 0.82 (compared to the prime aged), and for those above 60 it fell from 0.73 to 0.55. The result is clearly illustrated in table 4 through the comparisons of the age composition in the complete cohorts and the age composition in the survivor group. While the 50-59 group increased their share among the survivors (during the first two years) with 8.5 per cent (from 6.55 to 7.11) in the 1990-cohort, they increased their share with 94.6 per cent (from 8.37 to 16.29) in the 1996-cohort. The above 60 group increased their share with 236.8 per cent (from 2.88 to 9.70) in the 1990-cohort and 508.7 per cent (from 2.98 to 18.14) in the 1996-cohort.

For both cohorts, it is the case that access to unemployment benefits (or labour market program payments) reduces the exit probability. The effect seems to be have been stronger in 1990 (24 per cent reduction of the hazard) than in 1996 (13 per cent reduction of the hazard). Members of the 1990-cohort were eligible for unemployment benefits for a maximum duration of 80 weeks, and there seems to be a slight increase in the hazard rate as this duration is approached[15]. However, the cut-off period was limited to 12 weeks, after which a new period with benefits could begin (on a slightly lower replacement ratio). And during the cut-off period, many unemployed received social security contributions. For the 1996-cohort, there was in reality no limit to the duration of benefits.

Participation in labour market programs reduces the hazard during the program period (with 45 pr cent in 1990 and 51 per cent in 1996), while it raises the hazard as soon as the program is completed (by 36 per cent in 1990 and 25 per cent in 1996). One should, however, be careful in interpreting these estimates as reflecting ‘treatment effects’, as there is an obvious endogeneity problem present. Participation in a labour market program can hardly be considered an exogenous event with respect to persons labour market prospects; hence selection mechanisms (both self-selection and administrative selection) are likely to affect the results.

Figure 11 plots the hazard rates for both cohorts, based on the 1990-means for all explanatory variables. Hence, the two estimated hazard profiles are calculated for identical individuals facing identical labour market conditions (as measured by the aggregate outflow rate). The 1996-hazard lies almost everywhere above the 1990-hazard. Since the two hazard rates are estimated as a multiple of the two corresponding aggregate hazard rate, this suggest that newly unemployed workers did better compared to the existing pool of unemployed persons in 1996 than in 1990. A corollary is that the stock of unemployed workers in 1996 was less competitive (compared to newly unemployed workers) than their 1990-counterparts. Hence, the picture of an unemployment pool

15. Results in Røed et al, 1999, suggest that higher exits in the period around the benefit cut-off is more related to withdrawals from the labour forced than to higher transition rates to employment.

Figure 11.
Estimated Hazard rates for the January 1990 and the January 1996 unemployment cohorts.

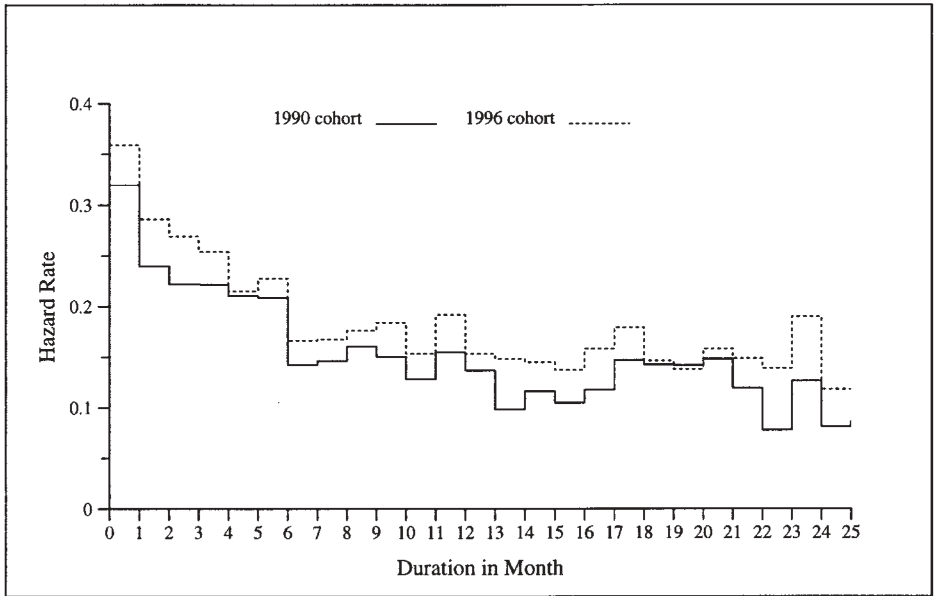

Note: The estimated hazard rates may be interpreted approximately as monthly probabilities of exiting unemployment, given that no exit has occurred so far. The hazard rates displayed are calculated for a the mean covariate vector in the 1990-cohort (to facilitate a direct comparison, this vector is also applied to the 1996-hazard).

consisting of relatively more hard-to-employs in 1996 than in 1990 is confirmed.

6 Concluding remarks

Our investigation into changes in the composition of unemployment leads us to draw the following conclusions:

1. There are no signs of low-education workers doing systematically worse than before, relative to high- or medium-education workers. Relative inflow rates have increased slightly, while relative outflow rates have decreased or are stable. In absolute terms, low-education workers are hit harder by economic slumps than high-education workers (i.e. their unemployment exposure increases more). In relative terms, there are no clear differences. The unemployment duration analysis comparing the cohorts from 1990 and 1996 indicates that low-education workers have improved their relative exit-probability slightly.
2. The unemployment performances of men and women have converged, both in terms of inflow, outflow and total exposure. The unemployment duration analysis also indicates that marriage now has significantly less negative impact than before on the exit probability of women.
3. The relative performances of different age groups have changed dramatically, and the change is a pure duration/exit phenomenon. While young persons exit faster than before, older workers seem to have become almost stuck in unemployment.
4. Total unemployment exposure has become more unequally distributed than before. A given yearly rate of unemployment is likely to be made up of fewer persons – each being unemployed more – now than what used to be the case. The average member of the unemployment pool has become less of a contestant in the competition for vacant jobs.

Even though we confirm the belief that little of interest has happened to the relative labour market performance of different educational groups, our results do suggest that some workers are systematically losing out to others. There are indications of increasing inequality in the distribution of the unemployment-burden. Hence, it is possible that some skill-biased changes in labour demand have occurred. But in that case, the relevant skill-measure is not educational attainment.

References

Andersen, A. S. and Aaberge, R. 1983, *Analyse av ulikhet i fordelingen av levekår.* Rapporter, 83/21, Statistisk sentralbyrå.

Bean, C. 1994, European Unemployment: A Survey. *Journal of Economic Literature*, Vol. XXXII, No. 2, 573-619.

Berg, H. and Børing, P. 1997, *Registrerte arbeidssøkere 1989–1995.* Rapport, 2, Arbeidsdirektoratet.

Blanchard, O. J and Diamond, P. 1994, Ranking, Unemployment Duration and Wages. *Review of Economic Studies*, Vol. 61, 417-434.

Burda, M. and Wyplosz, C. 1994, Gross Worker and Job Flows in Europe. *European Economic Review*, Vol. 38, No. 6, 1287-1315.

Butler, R. J. and McDonald, J. B. 1986, Trends in Unemployment Duration Data. *The Review of Economics and Statistics*, Vol. LXVIII, No. 4, 545-557.

Clark, K. B. and Summers, L. H. 1979, Labor Market Dynamics and Unemployment: A Reconsideration. *Brookings Papers on Economic Activity*, 13-60.

Cohen, D. Lefranc, A. and Saint-Paul, G. 1997, French unemployment: a transatlantic perspective. *Economic Policy*, Vol. 25, 265-291.

Hargreaves Heap, S. P. 1980, Choosing the Wrong Natural Rate: Accelerating Inflation or Decelerating Employment and Growth. *Economic Journal*, Vol. 90, 611-620.

Jackman, R. and Layard, R. 1991, Does Long-term Unemployment Reduce a Person's Chance of a Job? A Time-series Test. *Economica*, Vol. 58, No. 229, 93-106.

Jensen, A. M. and Jensen, P. 1997, Who Bears the Burden of Unemployment? Some Danish Evidence. Manuscript. Centre for Labour Market and Social Research, Aarhus.

Johansen, L. 1982, Some Notes on Employment and Unemployment with Heterogeneous Labour. *Nationaløkonomisk Tidsskrift*, special issue Economic essays in honour of Jørgen H. Gelting, København.

Krugman, P. 1994, Past and Prospective Causes of High Unemployment. Paper presented at the Jackson Hole Symposium, Wyoming, August 1994.

Machin, S. and Manning, A. 1998, The Causes and Consequences of Long-term Unemployment in Europe. Centre for Economic Performance, Discussion paper 400.

Makridakis, S., Wheelwright, S. C. and McGee, V. E. 1983, *Forecasting: Methods and Applications* (2nd ed.), John Wiley & Sons, New York.

Meyer, B. 1990, Unemployment Insurance and Unemployment Spells. *Econometrica*, Vol. 58, No. 4, 757-782.

Narendranathan, W. and Stewart, M. B. 1993, Modelling the Probability of Leaving Unemployment: Competing Risks Models with Flexible Base-line Hazards. *Applied Statistics*, Vol. 42, No. 1, 63-83.

Nickell S. and Bell, B. 1995, The Collapse in Demand for the Unskilled and Unemployment Across the OECD. *Oxford Review of Economic Policy*, Vol. 11, No. 1, 40-62.

Nickell S. and Bell, B. 1996, Distribution of Wages and Unemployment in OECD Countries. *American*

Economic Review, PAP, Vol. 86, No. 2, 302-308.

OECD 1985, Moving In and Out of Unemployment: The Incidence and Patterns of Recurrent Unemployment in Selected Countries. OECD Employment Outlook, 99-114.

OECD 1990, Employment Outlook. Paris.

OECD 1994, The Jobs Study. Paris.

Phelps, E. S. 1972, *Inflation Policy and Unemployment Theory: The Cost-Benefit Approach to Monetary Planning*. W.W. Norton, New York.

Pissarides, C. A. 1992, Loss of Skill During Unemployment and the Persistence of Employment Shocks. *Quarterly Journal of Economics*, Vol. 107, No. 4, 1371-1391.

Røed, K. Raaum, O. and Goldstein, H. 1999, Does Unemployment Cause Unemployment? Micro Evidence from Norway. *Applied Economics*, Vol. 31, No. 10, 1207-1218.

Viswanath, T. 1989, Job Search, Stigma Effect and Escape Rate from Unemployment. *Journal of Labor Economics*, Vol. 7, 487-502.