JAMES K. GALBRAITH
Professor of Economics, University of Texas at Austin, U.S.A

INEQUALITY AND UNEMPLOYMENT: AN ANALYSIS ACROSS TIME AND COUNTRIES

UDK 330.1KEYNES(063)=20

I. INTRODUCTION AND MAJOR FINDINGS

What is the relationship between inequality and unemployment? Many economists believe that these two variables move in opposite directions. They argue that flexible labor markets, characterized by a wider dispersion of earnings, will clear more readily and so yield lower rates of unemployment; correspondingly movements toward more flexibility should accompany reductions in unemployment. Loose comparisons between Europe, where wage structures are said to be compressed and unemployment is high, and the United States, where the reverse is true, have bolstered this belief over the past decade..

Very recently, though, doubts have emerged. In a recent study, the Organization for Economic Cooperation and Development reports that low-skilled workers are unlikely to face significantly higher rates of unemployment in countries with more equal wage distributions.[1] This finding was instantly controversial. A report in *The Economist*[2] challenged it, arguing that:

---

1 "Earnings Inequality, Low-Paid Employment and Earnings Mobility," *OECD Employment Outlook 1996.* Cited in *The Economist,* infra. Nickell and Bell (1996) report a finding similar to that of the OECD, for a study covering Germany, the UK and the United States.

2 "Jobs and wages revisited," *The Economist*, August 17, 1996, p 62.

"By itself, the lack of a positive correlation between high levels of inequality and low unemployment proves nothing. First, the analysis includes only a small number of countries... Second, and most important, the OECD analysis failed to take into account other factors that influence levels of unemployment, such as standards of education and training. Because these vary across countries, you would not expect to find a perfect ŠNB: presumably, inverseĆ correlation between levels of unemployment and levels of inequality at any single point of time."

"A more direct test of the value of flexible labour markets is how they respond to shifts over time... A comparison of changes in unemployment and changes in inequality... over the past decade and a half show a much closer correlation. Countries that have allowed the relative wages of low-skilled workers to fall have, in general, seen the smallest increases in unemployment."

The Economist supports this conclusion with a scatter chart covering ten countries over a single time frame, 1980 to 1995. The emphasis on change over time is surely both correct and useful. But, as I will show in this paper, everything else about the position quoted is completely wrong, and even the cautious OECD study finding "no correlation" between unemployment and relative equality understates the actual direction of the evidence.

This paper will show, for a group of countries mainly from Europe and North America[3] , that:

-- For most countries, the time-series relationship between unemployment and wage inequality is usually *positive* and in some cases strongly so. A few smaller countries have slightly compressed their wage structures while experiencing large increases in unemployment, resulting in a negative correlation. In no case has

[3] Countries covered are Austria, Belgium, Canada, Denmark, Finland, France, Germany, Italy, Netherlands, Norway, Spain (data begin in 1979), the UK and the US. Some parts of the analysis bring in information from Australia, Greece, Japan, Portugal, Sweden, and New Zealand.

any country systematically reduced unemployment while increasing inequality in the wage structure.

--- Year by year from 1970 through 1992, the correlation across countries between levels of unemployment and levels of inequality is overwhelmingly positive, not negative or even zero (as the OECD study apparently found). This correlation does start out negative, and stays so from 1970 through 1973. But after 1974 it is positive for every year up to the end of the data set, and the correlation coefficient is itself associated with the overall level of unemployment in the OECD. That is, higher average unemployment produces a stronger positive association between unemployment and inequality across countries.

--- Year by year, the correlation across countries between the *change* of unemployment and the *change* of inequality is sometimes positive and sometimes negative. The average value of the correlation coefficient across all years is virtually zero, with a high standard deviation. Thus a rising level of unemployment can accompany a rising or falling degree of inequality in the wage structure. There is no consistent pattern, except that when overall unemployment is high, the relationship between changes is more likely to be positive.

--- There is a complex pattern of international interdependence both in wage structures and unemployment rates, and these patterns of interdependence are only partly overlapping. There are strong elements of geographic affinity in both patterns.

Section II of this paper discusses some issues of theory, data, and method, and presents a benchmark estimate of inequality in the wage structure for the United States.

Section III presents year-by-year time-series measurements of inequality in the manufacturing wage structure for twenty-two OECD countries. These measures are self-consistent within countries and across time, but they are not comparable across countries. The section then shows how for thirteen of these countries the measurements can be converted to estimates of the change in the overall inequality of family income, measured by the Gini coefficient, by benchmarking to available Gini estimates from the Luxembourg Income Surveys. Thus it is possible to reverse-engineer annual estimates of earnings inequality, comparable across countries, for a

wide variety of countries going back to the early 1970s in most cases.

Section IV presents an analysis of the correlation between unemployment and inequality, with the main conclusions as outlined above. Section V presents an analysis of internatonal patterns of interdependence. Section VI offers conclusions and avenues for further research.

II. THEORY, DATA, METHOD

Strictly speaking, the neoclassical theory of income distribution does not predict any association between wage dispersion and unemployment. The theory merely says that compensation will follow the marginal productivity of labor. In the purest setting, there will be one single wage rate (the law of one price, in effect); in more realistic models wages may vary with risk, skill, compensating differentials and other factors. A systematic association between inequality and employment will emerge only if there exists some natural common distribution of underlying wage determinants across countries, and if some countries "unnaturally" compress their wage structures while others do not. This seems to be a widely-held view of the contrasting experience of Europe and the United States in recent years.

Theories derived from Keynes have a different implication. In the Keynesian case, rising unemployment may *cause* an increase in the dispersion of hourly wages, simply by breaking down the common inflexibility of relative nominal wage rates that characterizes periods of full employment. Keynes (1936) emphasizes the importance of relative pay differentials to workers. Surely it is an easy implication that some workers -- those strongly unionized, for instance -- defend themselves more effectively than others.

Thus a finding of a consistent inverse association across countries or through time between unemployment and inequality would support the prevailing pseudo-neo-classical view. As *The Economist* argued, this would be sufficient but not necessary. An

inverse correlation between rates of change of inequality and of unemployment would also support the neoclassical case. A finding of no association, either as regards levels or rates of change, would be inconclusive. Perhaps some countries are "naturally" more equal than others; perhaps those visibly compressing their wage structures are merely reflecting in law a more equal underlying distribution of skills and risks. On the other hand, a finding of *positive* association between hourly wage inequality and unemployment, or between changes in these variables, either across countries or through time, would be hard to square with the neoclassical vision. Such a finding would support the alternative Keynesian view.

The next question is, what to measure? Hourly average wages are the pay variable of theoretical interest; the theoretical issue is the relationship between unemployment and an hour's wage. But most data sets dealing with inequality rest not on the structure of individual hourly wages but on the distribution of family *incomes*. Factors affecting family income, such as changes in the demographic composition of families, have to be sorted through quite carefully before a measure based on incomes can be taken to be closely related to wages.

Another problem concerns the measure of inequality employed. Partial measures, such as a ratio of earnings at the lowest decile to those at the median -- a measure used in the OECD study -- can be unrepresentative of the structure of wages as a whole. Unemployment, after all, occurs throughout the structure of wages; the phenomena of downsizing and technological obsolescence can hit the highly placed and the highly skilled. Equally, changes in the minimum wage will be more important to a low-to-median ratio than they may be to the entire distribution. Ideally, then, one should seek a summary measure of inequality that covers the entire distribution of wages, based on individuals rather than families, on hours worked or a close approximation thereof, and if possible excluding non-wage sources of income.

We propose a direct measure of the change in inequality in wage structures that nearly meets these criteria. Our measure is

based on the work of Henri Theil,[4] whose measure of dispersion, derived from theories of information and entropy, is known as Theil's T. T has the very useful property of decomposability. For any set of mutually exclusive groups, T is the sum of the inequality between groups and the inequality within them. Thus a lower-bound estimate for T, T' or the between-group component of T, can be computed from grouped data. If grouped data are collected on a sufficiently detailed and consistent basis over time, then change in annual T' can serve as a time-series estimate of changes in inequality.

This insight greatly broadens the potential sources of data from which measures of inequality can be computed. In particular, it makes possible the use of industrial data sets, with measures of average annual earnings per employee by industry, for this purpose. Industrial classifications are, after all, merely sets of mutually exclusive groups. Their major limitations for this purpose is two-fold: (1) such data sets may have some sources of within-group inequality that do not vary between groups, and (2) the quality of industrial data is much better for manufacturing than for services; many such data sets do not cover the services sector at all.

We have examined both of these issues very extensively as regards the United States, using three alternative measurements of cross-industrial wage dispersion. The first is based on the Annual Survey of Manufactures for the years 1958-1992; in our analysis we reduced 139 three-digit industries into 23 internally-homogeneous groups through methods of numerical taxonomy beyond the scope of this paper.[5] Hourly wage data are used in this analysis, and employment weights are from the same source. The second is the OECD's Structural Analysis Database (STAN), which at its finest level of disaggregation breaks the industrial structure into about 40

---

4 Henri Theil, *Statistical Decomposition Analysis: With Applications in the Social and Administrative Sciences*, Amsterdam-London: North Holland Publishing Company,1972. The formula for T' is $T' = \Sigma(p_i \ \mu_i \ /\mu)\log(\mu_i \ / \ \mu)$, where the p_is represent the relative employment of the ith group, and $\mu_i \ /\mu$ is the ratio of the average hourly wage in the ith group to the global mean.

5 See Galbraith and Calmon, 1990, 1994, 1996 for a general description of methods, and Galbraith, forthcoming, for a detailed explanation of this research.

mutually exclusive groups, with the useful property that the same classification scheme is employed across all covered OECD members. STAN data on annual earnings per employee compiled in this way are now available for the years 1970 through 1992, with annual updates. The third U.S. data set was compiled for the years 1920 to 1947 by Thomas Ferguson and myself from data originally collected by the National Industrial Conference Board and other sources; it covers a wide range of services as well as manufacturing, and a full analysis of that material will be presented in a separate paper (Ferguson and Galbraith, forthcoming).

These analyses show that our postwar measures of T' based on hourly wages or annual earnings are very good proxies for each other. They also closely mirror changes in inequality more broadly measured, for example by the Current Population Survey's annual estimate of a Gini coefficient for family incomes in the United States. The Bureau of Labor Statistics has produced an annual series of the Gini coefficient, based on quintiles of the Current Population Survey of family incomes and going back to 1947. For the period 1958 to 1992, the correlation coefficient between this series and T' estimated from the ASM is 0.86. For 1970 to 1992, the correlation between T' estimated from the STAN database and the CPS Gini is 0.92. The correlations indicate that both series provide high-quality approximations of measurements otherwise obtainable only from expensive, and for many countries non-existent, annual micro-data.[6]

Why is this so? The reason apparently is that with a sufficiently fine disaggregation of industries, most sources of variation within industries also show up as between-group variation between some industries and some others. Just as the shape and movement of a fishnet approximates the shape and movement of the fish within, the dispersion of wages across the industrial classification structure covaries in a reasonably faithful way with the income distribution. And, we have discovered, many of the forces that affect the manufacturing wage structure also affect the dispersion of wages between manufacturing and services. This is partly because

---

6 Further issues and qualifications concerning Theil-based inequality measures and their use to make cross-country-comparable level estimates of overall inequality in incomes are discussed below and in the appendix.

many activities classed as services are closely linked to manufacturing, and partly because the economy is unitary: the same large forces affect all sectors, however they may be classified by government accountants.

Using the first of these correlations to plug the gap in our data for T' for the years 1947-1958, and benchmarking the 1920-1947 series to the estimated 1947 values, I have computed a complete series for wage inequality in the United States from 1920 to 1992. This series is presented as Figure 1, alongside a measure of the unemployment rate. The correspondence of the two series is very strong, and in fact some 55 percent of the variation in T' over the entire period is explained just by the movement of unemployment. This figure constitutes the first exhibit in our case that unemployment and inequality are positively associated.

Figure 1 **Inequality and Unemployment**
United States, 1920-1992

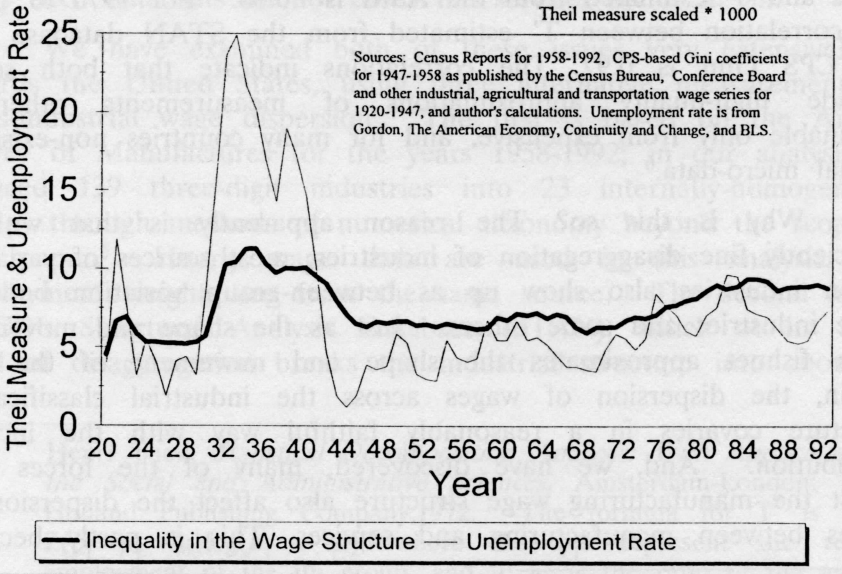

III. MEASUREMENTS OF WAGE INEQUALITY IN THE OECD

Figure 2 presents annual time-series estimates of the inter-industry dispersion of annual earnings for all the countries represented in the STAN database, for all years with sufficient usable data. This is the fundamental data set of the present study.

There are some problems with the data. Most notably, the sharp drop shown in wage inequality in Belgium in the early 1970s does not appear plausible, and does not show up when the same analysis is run at a higher level of industrial aggregation (eg, a two-digit decomposition). We therefore are inclined to regard that discontinuity as a data defect. A less visible defect apparently produces discontinuity in the data for France in 1977.[7] Still, it is remarkable how few such defects hit the eye.

[7] This defect emerged from a comparison of and apparent inconsistency between fixed and variable-weighted Theil measures, discussed below.

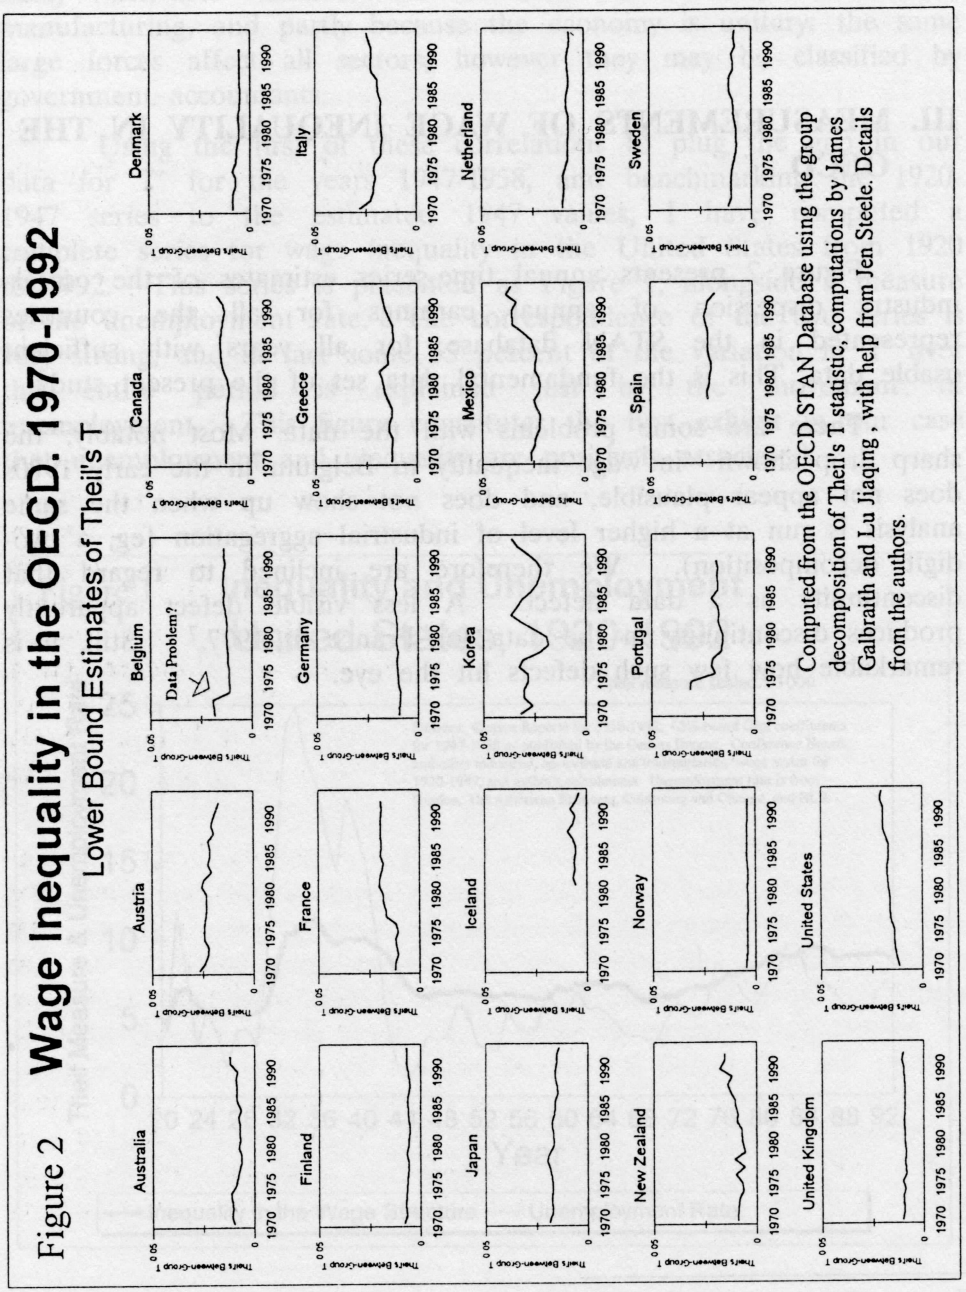
Figure 2 Wage Inequality in the OECD. 1970-1992
Lower Bound Estimates of Theil's T
Australia
Austria
Belgium
Data Problem?
Canada
Denmark
Finland
France
Germany
Greece
Italy
Japan
Iceland
Korea
Mexico
Netherland
New Zealand
Norway
Portugal
Spain
Sweden
United Kingdom
United States
Theil's Between-Group T
0.05
0
1970 1975 1980 1985 1990
Computed from the OECD STAN Database using the group decomposition of Theil's T statistic, computations by James Galbraith and Lu Jiaqing , with help from Jen Steele. Details from the authors.

Overall, the series appear broadly in accord with the stylized facts of the time and place. We observe that increases in inequality in Northern Europe appear quite low; in solidaristic Scandinavia wage structures have been almost undisturbed. In France in the 1970s, a known time of turmoil, inequality increases, as it does in Greece. In Canada, the U.S., Mexico and New Zealand there are sharp increases in inequality, especially toward the end of the period. The same is somewhat true of Italy. In Portugal following the revolution in 1974, inequality declined, the same happened in Spain under the Socialist government a decade or more later. A very few countries, notably Denmark and Austria, show steadily declining wage dispersion over the period.

Again, the level values of these series are not strictly comparable across countries. A finding that country A has a higher value of T' than country B does not necessarily imply that overall income inequality was greater in country A. Non-comparabilities may be due to cross-country differences in the composition of industrial employment, to cross-country differences in the proportion of total employment in covered manufacturing, and to cross-country differences in the allocation of non-wage incomes. However, given the the very high correlation between T' and the movement of inequality in family incomes observed in the United States, we do expect that *movements* in T' will prove a good proxy for the trend of inequality in most countries.

Fortunately, we have benchmark measures of inequality that are, in fact, designed to be comparable across countries. These are the Gini coefficients of the Luxembourg Income Studies, now available for scattered years for at least thirteen countries of the OECD's twenty-two. It is therefore possible to benchmark T' to the known level of the Gini for each of these countries. The simple assumption that the two series co-vary closely, known to be correct for the United States, yields a kind of reverse-engineered annual Gini coefficient for the thirteen countries, which we can call a Theil-Gini or T'G measure of inequality. Estimates of the Theil-Gini coefficient are presented in Figure 3, and a table of values is given in the appendix.

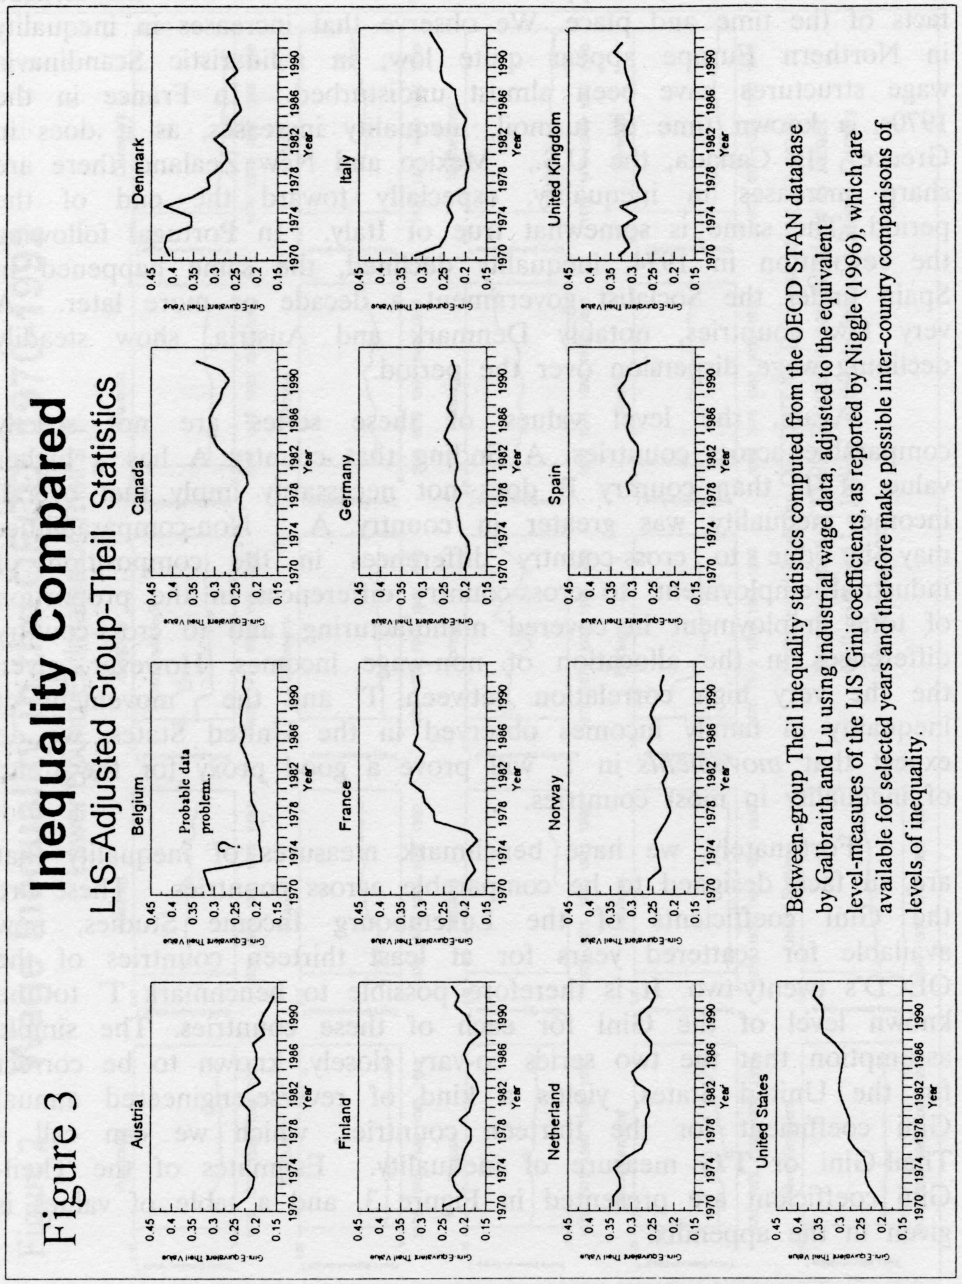

Between-group Theil inequality statistics computed from the OECD STAN database by Galbraith and Lu using industrial wage data, adjusted to the equivalent level-measures of the LIS Gini coefficients, as reported by Niggle (1996), which are available for selected years and therefore make possible inter-country comparisons of levels of inequality.

As a final exercise in this vein, we were tempted to exploit the additivity of the Theil series to produce an approximate meta series for changing inequality in Europe, incorporating both within-country and between-country components of wage inequality and adjusted to Gini-equivalent values. However, this effort did not succeed; the issues involved are also discussed in an appendix.

IV. UNEMPLOYMENT AND INEQUALITY

To return to the opening question of this essay, what is the relationship between inequality and unemployment? Figure 4 presents some evidence in the form of representative time trends for both series in eight OECD countries. In each case, unemployment is the thin line, inequality is the darker one; the inequality coefficient is scaled, usually by factor of 1000, to facilitate visual comparison. The patterns obviously vary. In some cases -- Germany, France, Canada, Australia -- the correspondence between the two series appears to be positive. In the case of Japan, the series move in opposite directions, as wage dispersion declines while unemployment rises. In the case of the UK, the series bear no apparent relationship to each other. For the U.S. it is hard for the eye to tell.

Figure 4

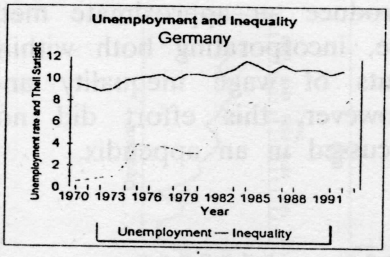

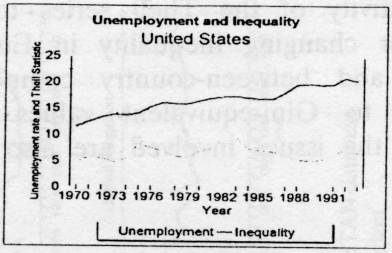

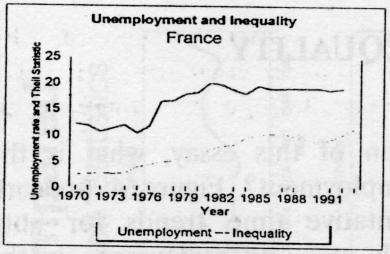

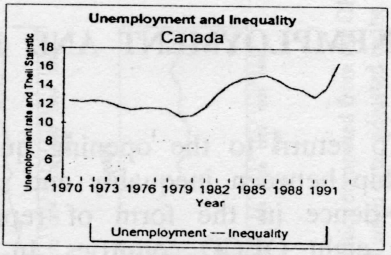

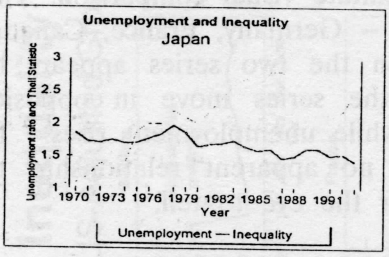

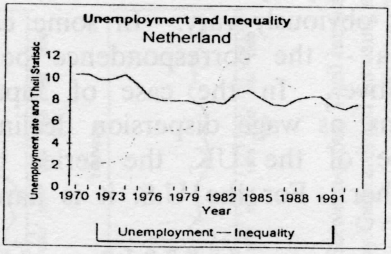

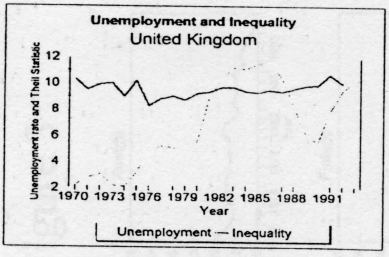

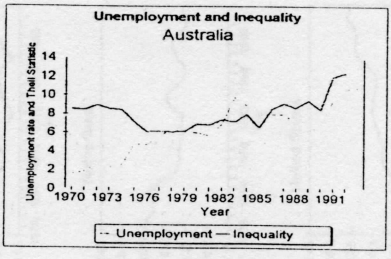

Unemployment and Inequality in Selected OECD Countries, 1970 - 1992

Figure 5 attempts to bring some order to this analysis. The figure presents an array of correlation coefficients between time-series measures of inequality and unemployment, ranked from lowest to highest across the eighteen countries for which internally-consistent time-series on both variables are available.

As the top half of the figure shows, Germany, France, New Zealand, Canada, Australia, the United States and Greece show a positive correlation between inequality and unemployment. In general, countries which experienced large increases in inequality, notably France, New Zealand, Canada, Australia and the United States, also experienced higher unemployment, and in general the movements occurred at the same time. The association is also strongly positive in the case of Germany, although the increase in wage dispersion in Germany is comparatively small.

Denmark, the Netherlands, Austria, Japan, Belgium, Finland, Norway and Sweden show a negative correlation between inequality and unemployment. For three countries, Italy, Spain and the United Kingdom, the correlation is essentially zero. The negative correlation is thus present mainly in those countries which succeeded in holding constant or even compressing their distribution of wages over this period.

The top half of Figure 5 is based on measures of inequality for which the underlying employment weights are allowed to vary from year to year according to actual shifts in the composition of employment. This permits accurate tracking of changes in wage dispersion in the manufacturing sector, but it introduces a bias into measures of the change in the structure of wages *per se*. For, as it turns out, there is a systematic association in many countries between changes in the composition of employment and changes in unemployment: as unemployment rises, people who are further away from (presumably, below) the mean of the wage structure are more likely to lose manufacturing jobs, and to be displaced into services or into unemployment. This source of bias is very small for the United States, but quite substantial in a number of other countries, particularly smaller ones.

The bias can be removed, yielding a measure of wage dispersion affected only by changing average within-industry wage rates, by fixing employment weights in the calculation of the

between-group Theil statistics at their 1970 values. The resulting fixed-weight inequality estimates are correlated to the unemployment rate in the bottom half of Figure 5. The results are striking. In eight countries, the correlation through time switches from *negative* to *positive*. In only one, Greece, the correlation moves the other way. In other words, when one considers only movements in relative wages, and not offsetting movements in the composition of manufacturing employment, the verdict that *rising* unemployment leads to *greater* inequality becomes stronger.

What accounts for the distribution of countries into positive and negative correlation cases and for the pattern of shifts from negative to positive correlations as one moves from variable to fixed employent weights? Several somewhat interconnected hypotheses suggest themselves. The economies with pronounced negative correlations between unemployment and inequality in the top half of Figure 5 and a tendency to flip to positive in the bottom half appear, on the whole, to be economies with a certain dualism in manufacturing, marked by strong and relatively stable, high-wage export sectors (Norway, Japan, Italy, Finland). In such countries, unemployment may rise more sharply among lower-wage manufacturing workers, even as their relative wages fall. This would diminish the weight of such workers in overall wage dispersion and account for the bias toward equality in the variable-weighted inequality measures for such countries. Large economies that sit, so to speak, at the center of their own economic basin (Germany, France, the United States) tend to show a positive correlation with unemployment on both inequality measures. So do the Anglo-Saxon economies (New Zealand, Canada, Australia, the United States). In these countries unemployment may be more nearly an equal-opportunity proposition.

Figure 5

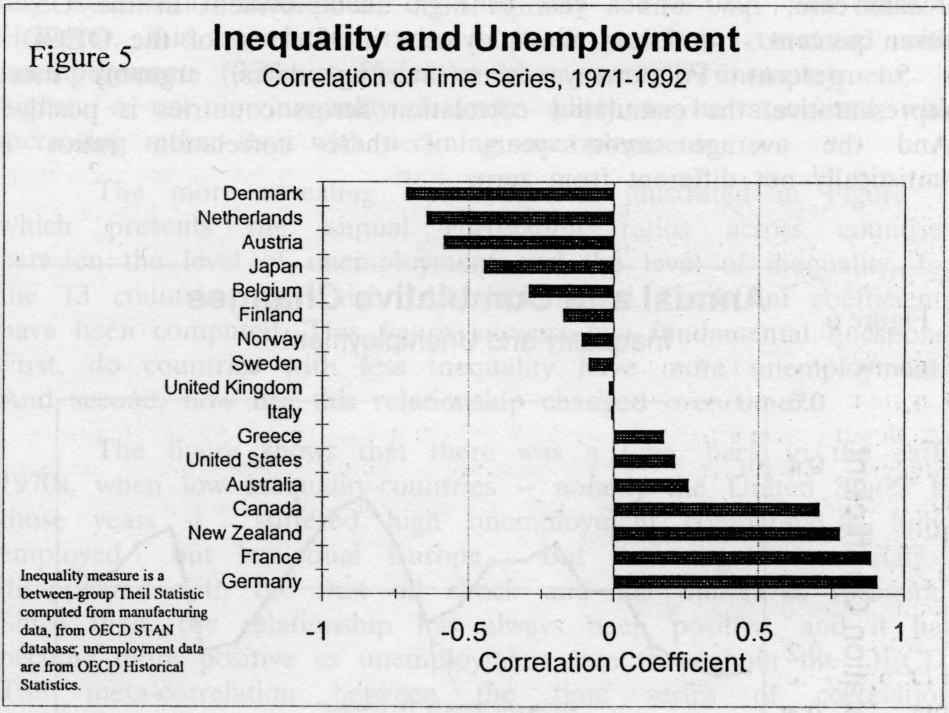

One might object that correlations across time of annual measures present a misleading picture, that one should instead examine the cumulative association of changes in inequality and changes in unemployment over a longer period. This is the procedure favored by *The Economist*, which presented data for ten countries showing the difference between changes in inequality and changes in unemployment beginning in 1980 and ending in 1995.

Such a procedure is however quite treacherous, for the correlation across countries of cumulative changes in unemployment and inequality turns out to depend critically on the choice of starting year. Figure 6 illustrates how the correlation coefficients vary, as one moves the starting dates forward from 1971 to 1991, with the ending date held constant at 1992. The comparison covers eighteen countries, including all ten in *The Economist's* scatterplot, and the results are consistent with what *The Economist* found for the dates it chose. For a starting date in 1980, the correlation over subsequent years between changing unemployment and changing inequality appears mildly negative, on the order of -0.25. But this is

a rare case; 1980 was a year of high unemployment in the US -- seven percent -- and low unemployment in the rest of the OECD -- 5.1 percent. For many other starting dates, arguably more representative, the cumulative correlation across countries is positive. And the average across years of these correlation ratios is statistically not different from zero.

Figure 6

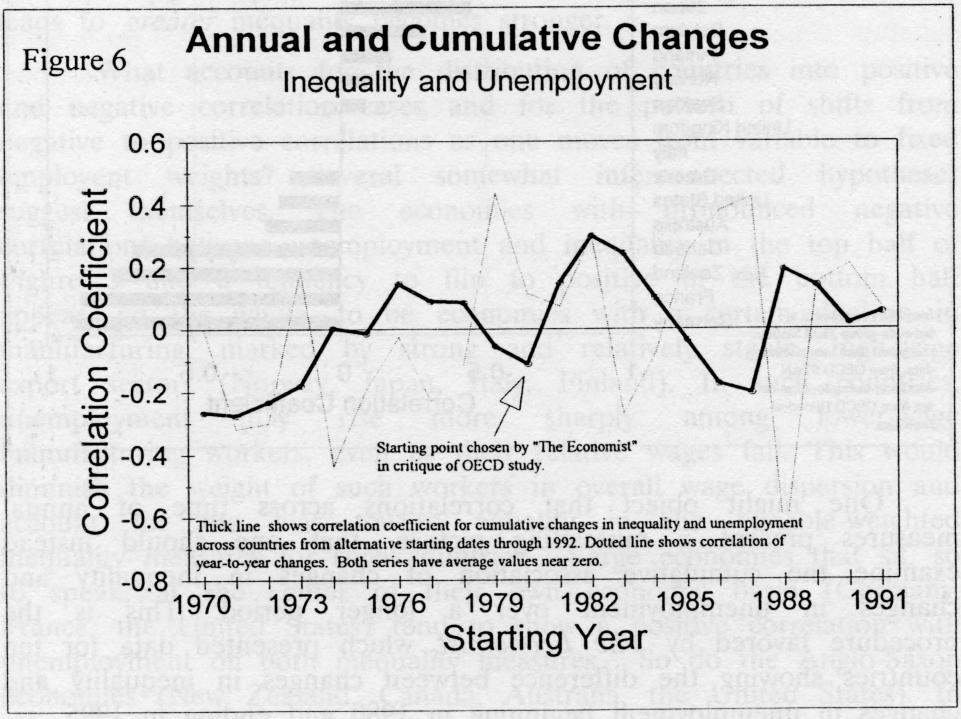

For completeness, I include a line on Figure 6 showing the relationship between the annual rate of change of inequality and the annual rate of change of unemployment, across countries, by year from 1971 to 1992. These correlation ratios of year-to-year change fluctuate erratically from one year to the next. There is no consistent pattern, and on average across years the correlation of these movements is zero. Unemployment can rise or fall, without

any systematic movement of inequality either up or down.[8] However, this correlation ratio is itself positively correlated with unemployment (0.23 coefficient): when overall unemployment is higher, a rise in inequality is more likely to be associated with increasing rather than with declining unemployment.

The more revealing procedure is illustrated in Figure 7, which presents the annual correlation ratios across countries between the level of unemployment and the level of inequality, for the 13 countries for which LIS-benchmarked Theil-Gini coefficients have been computed. This figure answers two fundamental questions. First, do countries with less inequality have more unemployment? And second, how has this relationship changed over time?

The figure shows that there was a time, back in the early 1970s, when low inequality-countries -- notably the United States in those years -- suffered high unemployment compared to fully-employed but less-equal Europe. But that negative association disappeared with the first oil shock and the mid-1970s recession. Since then, the relationship has always been positive, and it has become more positive as unemployment rose throughout the OECD. The meta-correlation between the time series of correlation coefficients (dark line) and the OECD unemployment rate (thin line) is .89, which seems high enough.

---

8 The same lack of clear association applies when fixed-weighted inequality measures are substituted. In that case twelve countries show a positive correlation, seven show a negative one, the average across countries is less that -0.02, twelve countries have correlation coefficients less than .15 in absolute value, and only one country shows a correlation of movements greater than 0.4 (absolute value).

Figure 7

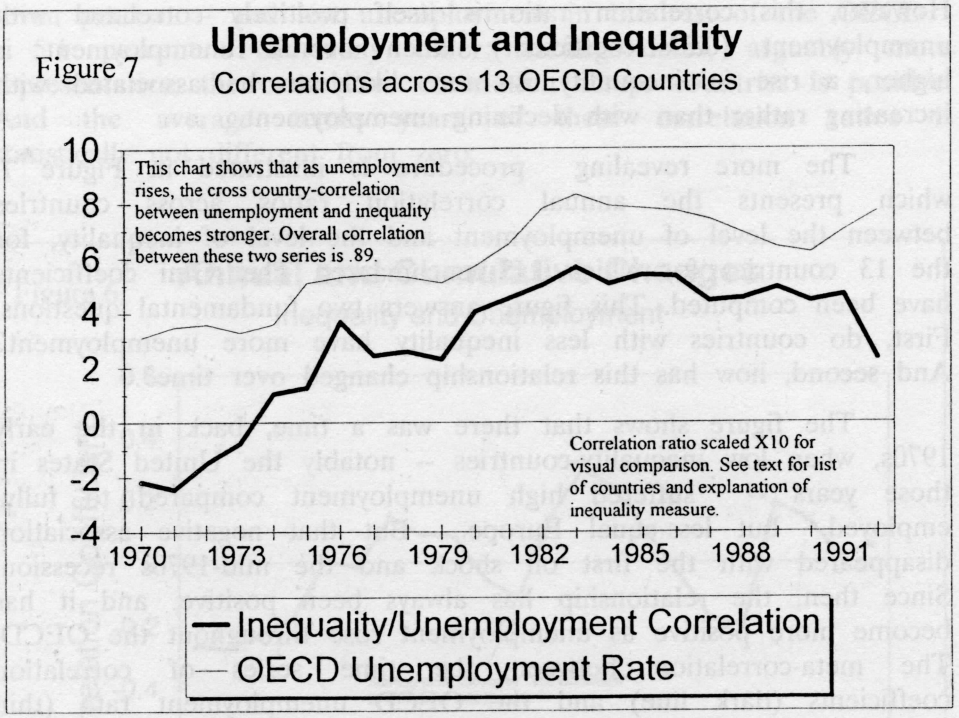

Thus we are at the situation foreshadowed in the introduction. There is a *positive* cross-country association between levels of unemployment and inequality. The within-country time-series correlation is positive for the larger, market-oriented countries, though negative for the corporatist cases. There is no meaningful cross-country association between movements of the two variables, either year-to-year or cumulatively.

All three findings are consistent with the Keynesian view, and difficult to reconcile with the pseudo-neoclassical vision. Ironically, an implication of these results is that the best way to restore the neoclassical worldview would be for the large, market-oriented and inequality-unemployment sensitive countries to embark on a program of global reflation. This would reduce both unemployment and inequality in the market-oriented countries, reducing the positive cross-country correlation between the two variables that presently exists.

V. INTERNATIONAL INTERDEPENDENCE

Our final exercise is a search for international patterns in the movement of inequality and unemployment. This search is facilitated by techniques of numerical taxonomy, specifically cluster analysis, where the criterion observation is the path through time, for each country, of the variable in question. In our case we use the annual rate of change of T', and the annual rate of change of unemployment. With either variable, the effect is to compose a taxonomy of nations, based on the closeness of their historical behavior as regards first inequality and then unemployment.

Specifically, we compute a matrix of Euclidean distances between vectors of rates-of-change, one for each country, and then group up our countries according the distance between them in (t-1) dimensional phase-space, where (t) is the number of years for which one has observations. A hierarchical agglomerative method that minimizes, at each step, the ratio of variance within groups to variance between groups (Ward's method), produces a tree diagram illustrating the covariation of each variable across the countries under analysis.

Figure 8 is such a tree diagram. Its geographical patterns emerge with great clarity. To be very precise, there is a North American pattern of (rising) inequality, to which the United States, Mexico and Canada all belong, and which has also influenced wage structures in Japan. Over on the right side of the diagram, one sees a similar North European orbit, including Austria, Norway and the Netherlands in close association with Germany, and, a little further out, Greece. Sweden and Finland form a Scandinavian pair in some association with the North European block; France is a bit more separate still. The UK, Italy and Denmark appear closer to the North American pattern of wage inequality. Finally, Korea, Australia and New Zealand each follow patterns of their own, not closely tied to any other country in this sample.

Figure 8

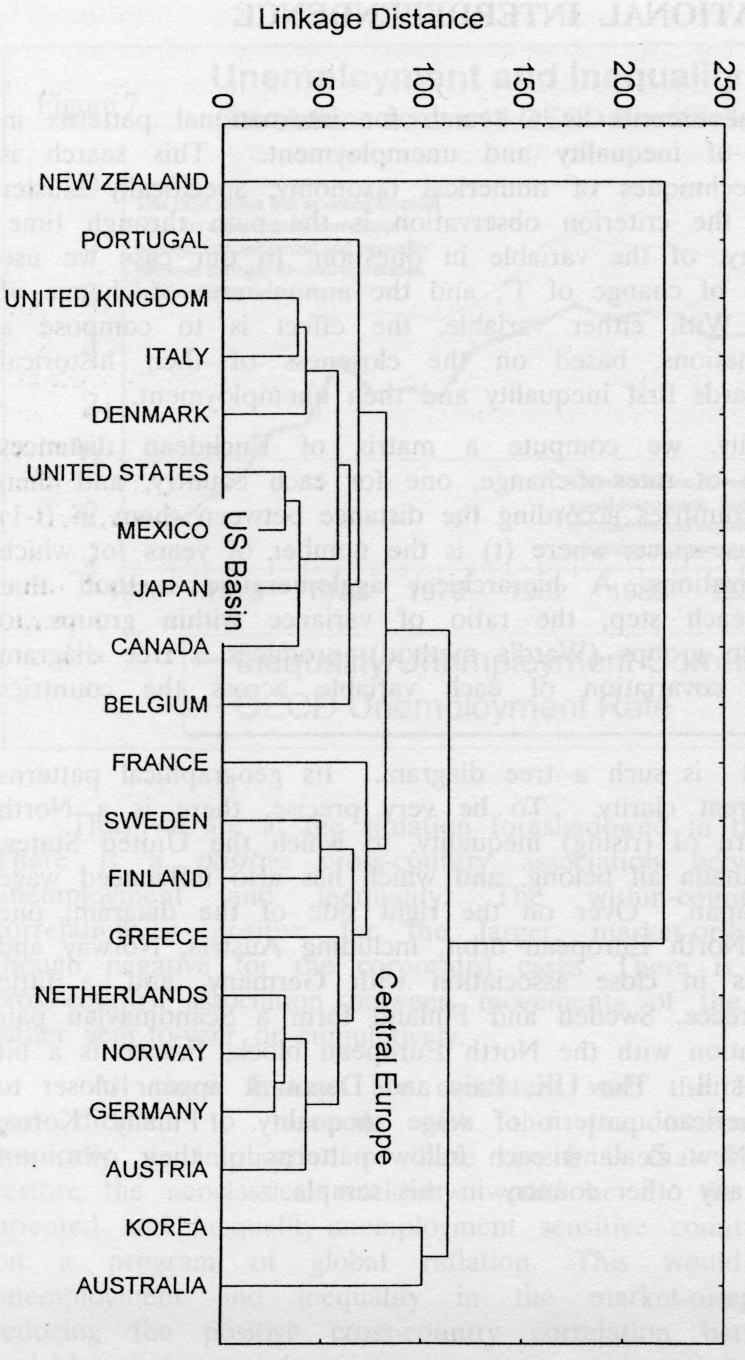

Patterns of Change in Wage Inequality

Cluster Analysis/Ward's Method/Euclidean Distances

Change in Between-Group Theil Entropy Measure of Inequality

OECD STAN Database, 1971-1992

Figure 8 suggests that the dispersion of wage structures is a transnational affair. It appears to be strongly influenced by geographical propinquity, trading patterns and perhaps other forms of transnational association. The existence of a Central European and of a North American inequality basin, so to speak, is an especially striking finding. Small countries apparently do not control their own fates, so far as wage structures are concerned.

Transnational patterns also characterize changes in unemployment, as Figure 9 illustrates. But the patterns are different. Once again there is a North American basin (unemployment data for Mexico were not readily at hand at the time this was written.) But now the UK joins the Netherlands, Germany, Sweden and Finland in a North European cluster. And there is a distinct South European cluster, including Italy, France, Belgium, Spain, Greece, and Austria, but also Japan. This cluster, however, has a behavior of unemployment rates which appears closer to the American orbit than to the North European one, and indeed more resembles the American pattern than the North European countries resemble each other. Norway, Portugal, Australia and Denmark are the outliers on this one.

Figure 9

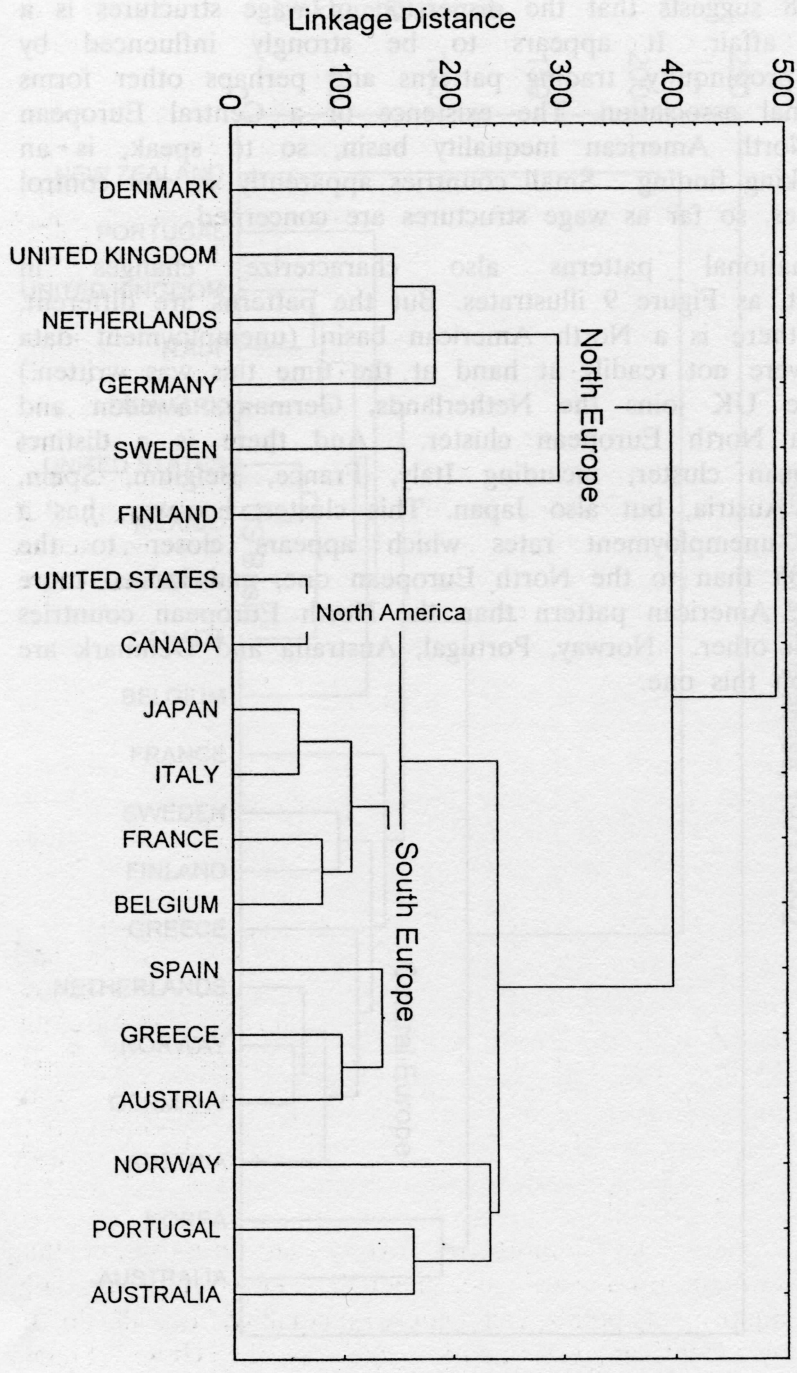

Patterns of Change in Unemployment

Cluster Analysis/Ward's Method/Euclidean Distances

OECD Historical Statistics 1961-1993

Overall, inter-country percentage fluctuations of unemployment, as measured by the linkage distance on the vertical axis, are about twice as large as inter-country percentage variations of inequality. And it would appear that wage structures and business cycles operate in overlapping basins. The wage structures of small countries appear to be quite tightly linked to their near neighbors, with Central European norms radiating out over much of the rest of Europe. But when it comes to the business cycle, it is the North American pattern that seems to rule most of the Western world.

Table 1 summarizes these findings in a rough three-by-three table. Changes in unemployment (across the top), and changes in inequality (down the side) are divided into three possible patterns: North American, North European (for unemployment) and Central European (for inequality), and Other. Countries are classed according to where they seem to best fit in the resulting 3x3 classification scheme. The result, while by no means definitive, seems a reasonable first cut at how patterns of international influence actually seem to run in the world.

VI. CONCLUSIONS AND FURTHER RESEARCH.

The most important conclusion of this paper is that it is possible and rewarding to exploit existing sources of wage and earnings data by industry to develop annual time-series measures of changing inequality, and that these measures do in fact usefully illuminate some of the central macroeconomic questions of our time. Correspondingly, studies that purport to draw deep conclusions from purely cross-sectional data, or from isolated point-to-point comparisons through time, or from small groups of countries are to be regarded with suspicion.

Second, the data support the Keynesian proposition that higher unemployment produces higher inequality, suggesting especially that a full employment policy will reduce inequality. It will do so within the larger continental economies, such as the United States,

and within the large European countries, such as France and Germany. There is no case anywhere in this evidence for the proposition that a country can reduce its unemployment rate by increasing the dispersion of its wage structure.

Third, this study raises the question whether there is anything sensible that a small country can do on its own about either its wage structure or its rate of unemployment. These variables in a small open economy appear to be substantially dependent on the winds blowing from larger countries, with the strongest winds on unemployment blowing from the direction of the United States. Thus, a coordinated global reflation led by America and seconded by Germany may be the only feasible path toward global full employment.

As for further research, there are many more countries with reasonable data on wages and employment across years and industries, and for whom computation of T' is therefore possible. As the Luxembourg Income Study expands its coverage, these estimates can be benchmarked and made comparable across countries. With Gini coefficients (or, better yet, sample-based Theil statistics) for multiple years, the ability of the Gini-adjusted Theil measure to track overall family income inequality through time can be checked. Further analysis of the determinants of wage inequality is warranted, taking account of the effects of changes in exchange rates, inflation rates, and other macroeconomic variables. In forthcoming work I will report extensively on the U.S. case; the project on wage structures at the University of Texas is presently working on other countries.

Table A1

Estimated Annual Values of the Theil-Gini Coefficient, 13 OECD Countries

	Austria	Belgium	Canada	Denmark	Finland	France	Germany	Italy	Netherlan	Norway	Spain	United Kingdom	United States
1970	0.248	0.311	0.255	0.408	0.251	0.200	0.203	0.318	0.342	0.261		0.331	0.200
1971	0.222	0.317	0.252	0.420	0.203	0.195	0.197	0.260	0.344	0.260		0.306	0.216
1972	0.208	0.325	0.256	0.414	0.222	0.176	0.189	0.245	0.326	0.231		0.319	0.231
1973	0.203	0.188	0.253	0.356	0.229	0.187	0.194	0.241	0.329	0.236		0.322	0.240
1974	0.224	0.191	0.243	0.356	0.211	0.197	0.210	0.229	0.342	0.240		0.289	0.238
1975	0.226	0.195	0.235	0.421	0.205	0.172	0.214	0.243	0.314	0.246		0.327	0.244
1976	0.224	0.193	0.240	0.353	0.198	0.193	0.218	0.233	0.275	0.233		0.267	0.266
1977	0.219	0.201	0.240	0.313	0.184	0.272	0.222	0.199	0.267	0.219		0.283	0.281
1978	0.215	0.199	0.238	0.318	0.184	0.275	0.211	0.200	0.265	0.208	0.308	0.290	0.278
1979	0.217	0.212	0.222	0.288	0.212	0.294	0.223	0.192	0.266	0.216	0.313	0.280	0.277
1980	0.215	0.214	0.225	0.271	0.188	0.301	0.223	0.197	0.258	0.222	0.295	0.297	0.287
1981	0.236	0.215	0.242	0.275	0.182	0.330	0.219	0.202	0.263	0.228	0.293	0.300	0.296
1982	0.228	0.213	0.272	0.301	0.168	0.325	0.225	0.202	0.257	0.217	0.286	0.310	0.301
1983	0.191	0.220	0.293	0.273	0.174	0.309	0.234	0.205	0.288	0.225	0.287	0.311	0.303
1984	0.208	0.226	0.305	0.261	0.193	0.296	0.250	0.197	0.299	0.246	0.304	0.301	0.298
1985	0.200	0.237	0.309	0.239	0.194	0.321	0.239	0.214	0.276	0.262	0.311	0.298	0.293
1986	0.200	0.230	0.313	0.247	0.202	0.312	0.229	0.221	0.256	0.245	0.306	0.304	0.293
1987	0.208	0.238	0.302	0.257	0.190	0.313	0.243	0.225	0.252	0.247	0.308	0.300	0.314
1988	0.182	0.237	0.287	0.274	0.164	0.315	0.232	0.254	0.273	0.257	0.321	0.308	0.340
1989	0.193	0.234	0.281	0.247	0.187	0.314	0.234	0.259	0.282	0.247	0.332	0.315	0.343
1990	0.184	0.224	0.266	0.260	0.199	0.315	0.234	0.255	0.282	0.227	0.315	0.317	0.340
1991	0.171	0.233	0.286	0.269	0.223	0.308	0.224	0.255	0.254	0.233	0.308	0.343	0.343
1992	0.159	0.230	0.342		0.214	0.313	0.231	0.269	0.244	0.226	0.259	0.322	0.364
1993					0.231			0.282	0.259				0.355

Estimates based on Theil lower-bound estimates of the change in inequality computed across industry groups, benchmarked to values of the Gini coefficient for family incomes reported by the Luxembourg income survey

Table A2

Year	European GT Estimate	United States GT	US CPS Gini
1970	0.265	0.200	0.353
1971	0.247	0.216	0.355
1972	0.241	0.231	0.359
1973	0.239	0.240	0.356
1974	0.236	0.238	0.355
1975	0.243	0.244	0.357
1976	0.230	0.266	0.358
1977	0.243	0.281	0.363
1978	0.247	0.278	0.363
1979	0.249	0.277	0.365
1980	0.251	0.287	0.365
1981	0.257	0.296	0.369
1982	0.261	0.301	0.380
1983	0.262	0.303	0.382
1984	0.264	0.298	0.383
1985	0.267	0.293	0.389
1986	0.264	0.293	0.392
1987	0.268	0.314	0.393
1988	0.272	0.340	0.395
1989	0.275	0.343	0.401
1990	0.271	0.340	0.396
1991	0.270	0.343	0.397
1992	0.270	0.364	0.403

REFERENCES

The Economist, "Jobs and wages revisited,", August 17, 1996, p 62.

Thomas Ferguson and James K. Galbraith, "The Wage Structure, 1920-1947." mimeo, The University of Texas at Austin, for the Harvard Economic History Seminar, March 8, 1996.

James K. Galbraith, *The Wage Structure and the Inequality Crisis*, in progress.

James K. Galbraith and Paulo Du Pin Calmon, "Wage Change and Trade Performance in U.S. Manufacturing Industries," Cambridge Journal of Economics, Volume 20, Nunber 4, July 1996, 433-450.

James K. Galbraith and Paulo Du Pin Calmon, "Industries, Trade and Wages," in Michael Bernstein and David Adler, Understanding American Economic Decline, New York: Cambridge University Press, 1994, 161-198.

James K. Galbraith and Paulo Du Pin Calmon, "Relative Wages and International Competitiveness," Lyndon B. Johnson School of Public Affairs Working Paper #56, January 1990.

W. Norton Grubb and Robert H. Wilson, "Trends in wage and salary inequality, 1967-1988," *Monthly Labor Review*, June 1992, 23-39.

John Maynard Keynes, *The General Theory of Employment Interest and Money,* London: MacMillan, 1936.

Stephen Nickell and Brian Bell, "Changes in the Distribution of Wages and Unemployment in OECD Countries, *American Economic Review Papers and Proceedings*, May 1996, 302-308.

Christopher Niggle, "New Evidence from the Luxembourg Income Surveys," mimeo, Redlands University, 1996.

Organization for Economic Cooperation and Development, "Earnings Inequality, Low-Paid Employment and Earnings Mobility," *OECD Economic Outlook 1996.*

Henri Theil, *Statistical Decomposition Analysis: With Applications in the Social and Administrative Sciences*, Amsterdam-London: North Holland Publishing Company,1972.

Edward N. Wolff, *Top Heavy: A study of the increasing inequality of wealth in Americam* New York: Twentieth Century Fund, 1994.

Appendix: The Theil measure and Gini-equivalent estimates

Originally drawn from information theory, Theil's T has the following formula:

$$T = (1/n)\ \Sigma(Y_i\ /\mu)\log(Y_i\ /\mu) \qquad (1)$$

Here, n is the number of individuals, Y_i is each person's income, and μ is average income for the whole population. "Log" is the natural logarithm.

Notice that, whenever a group population consists of equal individuals, the final terms in T all reduce to log $(y_i\ /\mu)$ = log(1), which is equal to zero. Thus T overall is zero for the case of perfect equality. And T increases, as deviations away from the average value increase. Since deviations of (y_i/μ) below the mean have values between zero and one, whereas deviations above the mean are unbounded, T increases as more of the observations move away from the average. Thus T is a reasonable way to measure the degree of dispersion about the average value for any group of observations, and that is, after all, what inequality is.

The formula for computing T from grouped data is this:

$$T = \Sigma(p_i\ \mu_i\ /\mu)\log(\mu_i\ /\ \mu) + \Sigma(p_i\ u_i\ /\ \mu)T_i \qquad (2)$$

where now p_i is the proportion of workers employed in the i-th group, μ_i represents the average income for the i-th group, μ represents average overall income, and T_i is the Theil T as measured strictly within the i-th group. Thus the grouped Theil statistic is the weighted sum of that part of inequality that occurs between groups (on the left of the above expression) and a part that occurs within groups (on the right).

The formula for T', the between-group-Theil statistic, is just the first element in the forumula for computing the Theil T from grouped data:

$$T' = \Sigma(p_i \mu_i / \mu)\log(\mu_i / \mu) \quad (3)$$

Since the within-group element in variation is omitted, this is obviously a lower-bound estimate of dispersion.

Our estimates of T' for individual OECD countries use average annual earnings in 40 manufacturing industries, weighted by employment in those industries. To compute a cross-country comparable Theil-Gini coefficient for a single country, the T'G index, we assume that T' is perfectly correlated to the unknown series of Gini coefficients (a fact known to be nearly true for the United States). We then benchmark the T' series to the known value of the Gini for whatever year it may be available from the Luxembourg Income Survey. To do this we simply scale up the whole T' series so that it matches the known Gini in the known year. The proportionately adjusted values of T' in other years give an estimate of what the Gini might have been, if a survey had been taken in those years.

We find that a series so constructed from the STAN for the United States is highly correlated with an actual measure of the Gini coefficient for variations in family income, based on the Current Population Survey, for the years 1970-1992 (rho = .92). However, several cautions are in order. First, our benchmark value for the U.S. Gini coefficient from the Luxembourg Income Survey, 0.343 in 1991, is well below the Census estimate of .397 for that year. Correspondingly, all other values in our series will be lower. We have not inquired into the sources of this discrepancy.

Second, our T'G series for the U.S. has about twice the standard deviation of the CPS Gini series. Our series increases much more sharply than the CPS series, and therefore projects back over twenty years to much lower estimated Gini values in the early 1970s than does the Census series. This turns the United States into a low-earnings-inequality country, relative to much of Europe, as recently as 25 years ago.

But we think this result is, at least, plausible in historical context. Inequality in U.S. wage incomes in 1970 would have been much lower than inequality in overall family incomes, due to the highly skewed distribution of non-wage incomes in the United States, a country with a very weak social welfare system, a high

concentration of private capital ownership and few nationalized industries when compared to Western Europe at that time.

It also seems likely that non-wage sources of income became more equal in the United States at least through 1980, because social insurance incomes rose during that time, and because the distribution of wealth tended to become somewhat more equal as inflation raised the value of houses relative to financial assets such as bonds whose ownership is highly concentrated (Wolff, 1994). Thus, U.S. wage and earnings inequality would have risen more rapidly than overall inequality in this period. The implication is that the dispersion of wage and salary earnings must have been quite low by international standards in the United States at the end of 1969. This was a time of full employment, strong labor unions, recently-installed trade protection in textiles (the multi-fiber agreements), and a shooting war, all of which work to compress the distribution of earnings.

Since changes in social insurance coverage and in the wealth distribution are likely to have been less important in Europe than in the United States in the decade of the 1970s, we think that the T'G estimates for Europe may be better reflectors of the trend in overall *family income* inequality than they are for the United States. However, for both regions, the T'G indicators should provide fairly robust estimates of the trend in inequality of wage and salary *earnings*, since the cross-industrial dispersion of such earnings is what they actually measure.

The absence of data on services from the STAN data set is another troubling issue. It is possible that a data set covering average wages in services would show a different rate of change of inequality than our restricted data set covering manufacturing. But would the rate be higher or lower? We can think of no clear *a priori* reason for bias in either direction. Some tentative evidence for the United States suggests that a significant part of wages in the services sector in a large-market economy covary with manufacturing wages, so that our data on manufacturing wages represent a larger fraction of total employment than they actually cover. In smaller countries, solidaristic wage structures are likely to extend across manufacturing and services, further reducing the likelihood that inequality measures taken on the manufacturing sector grossly misrepresent the larger picture. There is no reason

to think that overall inequality measures could run counter to the trend within manufacturing, and the American case suggests that this need not be a concern. The issue posed by the absence of services data is one of magnitudes, not direction.

To compute a Gini coefficient for Europe as a whole, we would need to start by deriving an estimate of T for Europe, from equation (2) above, with both between-country and within-country components. The p_i are available in our relative employment weights, and these can be treated as though they were approximate population weights. The u_i/u can be taken from ratios of within-country per capita GDP to a cross-country employment-weighted average. We computed these from OECD Historical Statistics using a1993 purchasing power parities.

Unfortunately, this process founders on the fact that T' is only a lower-bound estimate of within-country inequality, so that we have no good proxy for the full within-country T_i. Our best guess, noting that the values of T_{US} estimated for the United States by Grubb and Wilson are on the same order (about .35) as Gini coefficients, is that between-country variations in income have contributed fairly little (on the order of .02) to the overall inequality of European wage structures since 1970. These are all developed countries, and the relative incomes of the large countries are fairly close and change little, when PPP-adjusted. Ignoring the between-country variations, the table below provides a highly approximate estimate of trans-European inequality, basically a weighted average of within-country T'G estimates for eleven countries. This series is compared to the T'G estimate of earnings inequality in the United States.

If the table is approximately right, both societies are markedly less equal in their wage structures than was the United States, as recently as 1970, when the unemployment rate was last below four percent. But the United States has clearly lost its claim to being the middle-class, or classless society, in the intervening time. That distinction now belongs to Western Europe if it belongs anywhere at all.