

KEYNES LECTURE IN ECONOMICS

What is Labour-Market Flexibility? What is it Good for?

ROBERT M. SOLOW

*Massachusetts Institute of Technology
Fellow of the Academy*

THE EARLIEST KEYNES LECTURES tended to be on Keynesian subjects, either Keynes himself or the parts of economics on which he was a major influence. More recently that habit of piety has weakened or disappeared. I intend to go back to it, but not in the sense that I will be concerned with what Keynes or later Keynesians thought about this or that. The connection is more abstract. The macroeconomic role of the labour market is at the very centre of all that current talk about high and persistent European unemployment and its sharp contrast with the recent experience of the United States. The macroeconomic role of the labour market is also at the heart of one version—more or less the standard version—of Keynesian economics.

One legacy of *The General Theory*—much disputed—is the notion that the true source of excessive unemployment need not be some flaw in the labour market itself. In another standard version of the story, however, a lower nominal wage would correspond to higher employment, with the main chain of causation running through a lower price level, higher real money stock, lower interest rate, and higher aggregate real expenditure. In that story a higher nominal money stock would work just the same, without the need for deflation. From this angle, nevertheless, it looks as if nominal-wage-rigidity is the root of the problem. But other arguments, to be found in the *General Theory* itself, in Christopher Dow's Keynes Lecture of a few years ago, and in more recent work by Frank Hahn and myself, claim that full nominal wage flexibility would bring with it enough other problems to make it a non-solution to the

Proceedings of the British Academy, 97, 189–211. © The British Academy 1998.
Read at the Academy 30 October 1997.

problem of unemployment. It is certainly not progress to *define* complete wage-flexibility as the achievement at every instant of a nominal wage at which employment will be 'full' and then to declare that wage-flexibility is the solution to the problem of excessive unemployment.

All this will remind you of your youth, assuming that you are not young any more. It is not the trail that I want to follow. I propose to start from the other end. Measured unemployment has been very high in Europe for more than a decade, as compared both with the period before 1975 and with the United States currently. Almost uniformly, in public discussion of what is after all a high-visibility issue, the blame for this failure falls on 'rigidities in the labour market'. That response has already achieved the status of a reflex, exhibited by any central banker you might care to ask. The same knee-jerk reaction appears also as the conclusion of the OECD Jobs Study (though it might more appropriately be described as the assumption of the OECD Jobs Study). My plan is to start from that end and work backwards to something more precise and more analytical. It is only fair to say that there has been some serious good sense written on this topic; an excellent example is the 1995 pamphlet by G. Alogoskoufis, C. Bean, *et al.*, *Unemployment: Choices for Europe*, published by the Centre for Economic Policy Research in London. I want to mention also an interesting article by Stephen Nickell in the Summer 1997 *Journal of Economic Perspectives*.

My first observation is that 'labour-market rigidity' is never defined very precisely or directly in this context, but only by the enumeration of tell-tale symptoms. Thus a labour market is inflexible if the level of unemployment-insurance benefits is too high or their duration is too long, or if there are too many restrictions on the freedom of employers to fire and to hire, or if the permissible hours of work are too tightly regulated, or if excessively generous compensation for overtime work is mandated, or if trade unions have too much power to protect incumbent workers against competition and to control the flow of work at the site of production, or perhaps if statutory health and safety regulations are too stringent. It seems clear that those who point to labour-market rigidity as the source of high unemployment have something other than simple nominal or real wage rigidity in mind, or so I shall assume.

This sort of definition by example is far from satisfactory. Not that the examples are irrelevant: each of the restrictions I have mentioned certainly contributes its mite to labour-market rigidity in the very broad sense that it limits the possible responses to any exogenous change in circumstances. Nevertheless there are (at least) two important reasons to look for something more systematic.

Every one of these regulations or restrictions was intended to promote a desirable social purpose. Some may do so ineffectively or inefficiently. That is worth knowing; but the fact remains that wholesale elimination of these 'rigidities' is neither desirable nor feasible. They might be modified or traded

off against one another; but to do that intelligently already presupposes some sort of yardstick. One needs to form some idea of how much a particular regulation contributes to overall labour-market rigidity; in other words, one needs some roughly quantifiable measure of rigidity. Such a measure would serve other purposes as well: comparisons from time to time and place to place, for example, require a way of talking about the overall degree of rigidity.

The second reason is more parochial, but still important. If 'labour-market rigidity' is to be more than a slogan, it needs to be incorporated into macro-economic model-building, into the normal way we think about the determination of wages and employment in modern market economies. The catalogue of restrictions might be—and has been—studied one at a time from this point of view. But it is too hard to take account of them all at once. The analytical foundations of policy in this field would be strengthened if there were a useful summary indicator of labour-market rigidity that could function as an input into careful thinking about the macroeconomics of the labour market.

My candidate for this summary indicator is the location of the 'Beveridge curve' in whatever economy we are discussing. The Beveridge curve is the (negatively sloped) relation between the vacancy rate (the number of unfilled jobs expressed as a proportion of the labour force) and the unemployment rate (the number of unemployed job-seekers expressed as a proportion of the labour force). One can usually find such a curve in national data, if the data exist. It has negative slope for the commonsense reason that jobs are easier to fill, and the vacancy rate therefore lower, the more unemployed workers there are for employers to choose among. (This property can of course be derived from more primitive assumptions.) It is named after Sir William Beveridge, whose famous wartime report *Full Employment in a Free Society* defined 'full employment' to be a state of affairs in which the number of unfilled jobs equals the number of unemployed workers; this definition picks out the intersection of the Beveridge curve and a 45-degree line from the origin, when the unemployment rate is measured on one axis and the vacancy rate on the other.

A perfectly flexible labour market would then be one that interposes no obstacle to the frictionless matching of an unfilled job and an unemployed worker with the appropriate skills. In that case vacancies and unemployment could never coexist. The Beveridge curve would coincide with the axes of the diagram; there could be vacancies with no unemployment or there could be unemployment with no vacancies. Of course no real labour market could be perfectly flexible in that sense. Labour-market rigidities (including skill mismatches now) are precisely what allows vacancies and unemployment to coexist, and the more rigidities there are, the further the Beveridge curve

diverges from the limiting case, the further from the zero-zero point it is located.

One can test this interpretation against the list of symptoms I mentioned earlier. Generous unemployment insurance benefits allow unemployed workers to reject vacant jobs in the hope of finding something better; the ‘something better’ might include not working for a while. Restrictions on hiring obviously make it harder to match a vacancy with an unemployed worker. Restrictions on firing are more subtle; they may inhibit the creation of vacancies, but they also induce employers to be pickier because a bad match is costlier to undo. In fact anything that limits the employer’s control over performance is likely to shift the Beveridge curve adversely because adjustment to bad hires will be more difficult. Just thinking in terms of the Beveridge curve makes it easier to extend the list of effective labour-market rigidities. For instance, anything—statute, union rule, regulation or custom—that limits the geographical, industrial, or occupational mobility of workers, indeed anything that leads to segmentation of the labour market, is bound to shift the Beveridge curve adversely, because vacancies in one segment cannot be matched with unemployed workers in another. Obviously I have been ignoring genuine heterogeneity by tacitly including all workers and all jobs in one labour market. The real situation is more complicated but does not have to be dealt with in this sort of informal exposition.

Wage inflexibilities, including a high reservation wage, may also push the Beveridge curve away from the origin by preventing hires that could have been made if the wage were more flexible. I would like to leave it at that for now, because a complete theory of the labour market is not needed for the simpler point I am trying to make now: a good way to summarise the degree of rigidity in any national market is to see how far its Beveridge curve is from the limiting, unachievable case of perfect flexibility.

The theory of the Beveridge curve is in pretty good shape. (The best discussion is by Olivier Blanchard and Peter Diamond in ‘The Beveridge Curve’, *Brookings Papers on Economic Activity*, 1989, No. 1, pp. 1–60 and 74–6; they have a handful of other papers, including a brief introduction in ‘The Flow Approach to Labor Markets’, *American Economic Review (Papers and Proceedings)*, v. 82 (May 1992), 354–9.) The real problem is empirical. The very concept of a count of job vacancies is vague around the edges, and the same can be said of a count of unemployed workers. Employers can be more or less serious about filling a job, just as people can be more or less serious about finding a job. But national unemployment counts exist almost everywhere, whereas vacancy statistics are quite rare.

In the United States, official vacancy statistics were collected only for a brief interval in the 1960s. Instead the custom is to use a privately-collected time series of the volume of help-wanted advertising in newspapers. This is

obviously an imperfect surrogate; for example, a vacancy may be advertised several times or not at all, and in any case the role of the newspaper as an advertising medium has been changing. But it is a lot better than nothing. In Europe the availability of vacancy statistics differs from country to country, and there are occasional changes in definitions and methods. But there is enough to get on with.

The sort of labour-market model encouraged by the use of the Beveridge curve allows one to talk of labour-market flexibility as distinct from simple wage-flexibility. I have taken that opportunity for a couple of reasons. Those who talk about the need for more flexibility in European labour markets are presumably not just asking for more wage-flexibility; if they were, they could say so directly. Besides, the macroeconomics of wages is a very long story, going back at least to Pigou's *Theory of Unemployment*; there is nothing to be gained by bringing it up here. There is, however, one relevant empirical point that I would like to underline.

There was a time, in the early and middle 1980s, when the 'wage-gap hypothesis' was a leading candidate explanation for what was even then seen as unusually high unemployment for Europe. (The main reference is *The Economics of Worldwide Stagflation*, 1985, by Michael Bruno and Jeffrey Sachs.) The hypothesis was that real wages in Europe had outrun labour productivity. Among the consequences were low profitability, low investment, and a lot of unemployment. It is easy to see how this could have happened. Real wages typically move with labour productivity. An unexpected productivity slowdown began some time around 1970 (and continues still, though no longer unexpected); slow adaptation to this change could account for the opening of a wage-gap if real wages continued to reflect inertia induced by older expectations. One common extension was the idea that in Europe real wages were sticky; in the US nominal wages were sticky so the real wage could be 'inflated down'.

To say that the real wage has outrun the productivity of labour is to say that the share of wages in aggregate output has risen; and in fact the profit share in the major Continental economies was unusually low from about 1975 to the early 1980s.

Beginning in the early 1980s, however, there was a remarkable distributional shift to profits. The wage share in Europe began to fall, and may not yet have stopped falling. By now the wage share on the Continent is substantially lower than in North America. The wage-gap has disappeared, more than disappeared so to say, but the unemployment lingers on. The significance of this fact is that one cannot build a really convincing story about current unemployment that rests primarily on wage-rigidity that holds the economy-wide real wage at too high a level. The dynamics do not fit.

It is worth a reminder that the unemployment rate in the UK climbed to

very high levels in the years 1981–7 and has since receded, not to where it was in the 1970s but to a figure substantially lower than in France and Germany. During this period the profit share fluctuated around an essentially horizontal trend; here too the profit share has been rising in the 1990s, but it is still ever so slightly lower than in the mid-1980s. There is no particular comfort for the wage-gap story here either.

The Beveridge curves provide a somewhat heterodox perspective on the role of labour-market rigidities. (Here and elsewhere I am deeply indebted to Professor James Medoff of Harvard for his pioneering empirical work and his generosity in helping out with data and analysis. The data plotted in the European Beveridge-curve diagrams are all extracted from official sources. The figures are not really comparable from country to country, not even the unemployment rates. They are useful primarily for comparisons over time for each country separately.) I begin with the US (Fig. 1) because it is the natural benchmark for comparison with France, Germany, and the UK.

The diagram plots the ‘vacancy rate’ vertically and the unemployment rate horizontally, on an annual basis. Remember that the ‘vacancy rate’ is really an index of help-wanted advertising normalised by the labour force. It would be possible to doctor the data: the unemployment rate could be corrected for demographic changes, as George Perry suggested long ago; and the help-wanted index could be keyed to the brief period when vacancies were actually measured, and could be adjusted for the change in the importance of newspapers as an advertising medium, as Katharine Abraham showed a decade ago. I have omitted such refinements because it is only the qualitative picture that matters for now.

That qualitative picture stands out clearly. From 1958 through 1971, the US seemed to move along a well-defined Beveridge curve. During 1972–4 the curve shifted adversely, and settled for 1975–86 about three percentage points of unemployment to the right of its earlier location. Then, in 1987 and 1988, the curve seemed to return to its initial position, and has remained there for the past decade. One can interpret this as saying that the US experienced an episode of acute labour-market rigidity between the early 1970s and the early 1980s, and has now reverted to form. (Blanchard and Diamond, in the 1989 paper already cited, produce a monthly Beveridge curve for a different period, 1952–88, using slightly different data. The general evolution is entirely consistent with what I have just described.) Does this make sense?

I can invent a libretto to go along with that melody. The story line includes the pronounced productivity slowdown, leading first to an unrealistic reservation wage, and then to an eventual adjustment to reality. One could also make something of increasing segmentation of the labour market as older manufacturing industries decayed, and the economic structure shifted in favour of the service sector and the Sunbelt. I called this a libretto precisely to underline the

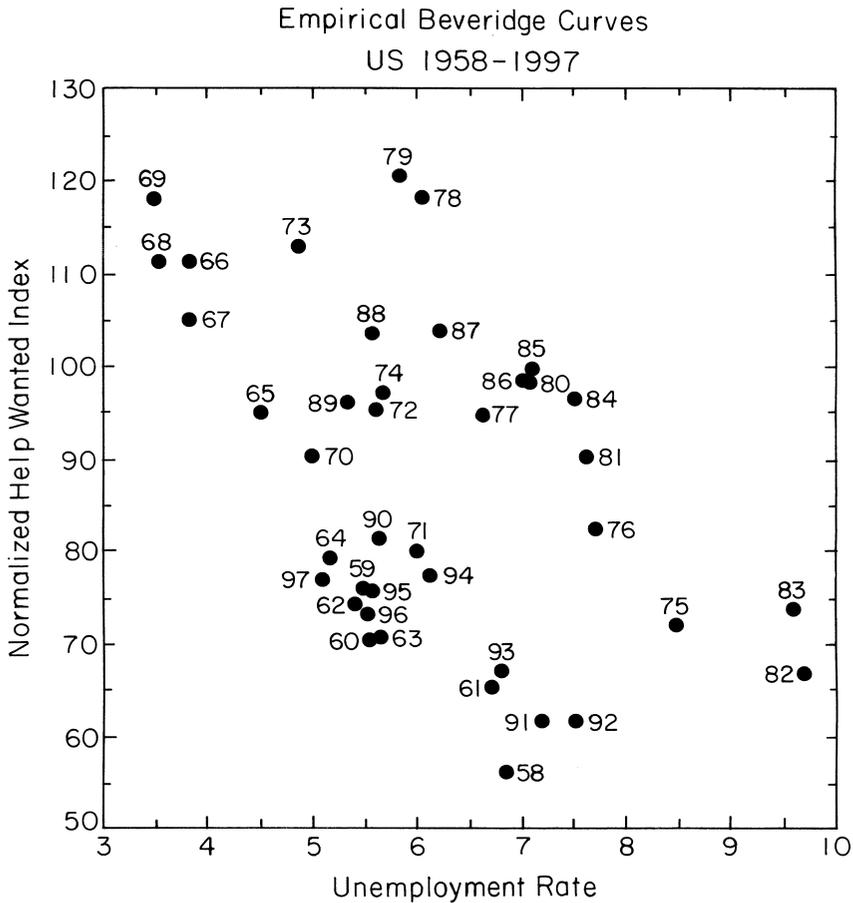


Figure 1.

lightweight character of the exercise. One can always invent a plausible story to cover a single episode; in this case the episode lasted for forty years, at least, and had three acts. It is worth noting, however, that the sorts of scenery emphasised in the usual version of the European opera do not seem to have been on stage in the US.

The picture in the UK (Fig. 2) as shown in the second diagram, is more complicated. Perhaps bemused by what happened in the US, I am inclined to push my luck and suggest an analogous, though not quite similar, evolution here. A determined reader of tea leaves could certainly see an initial Beveridge curve for the years 1964-72. Beginning in 1973, during a period of mostly rising unemployment, the whole curve seems to migrate to the right, settling down from 1983 to 1987, and then moving leftward again to what looks like a

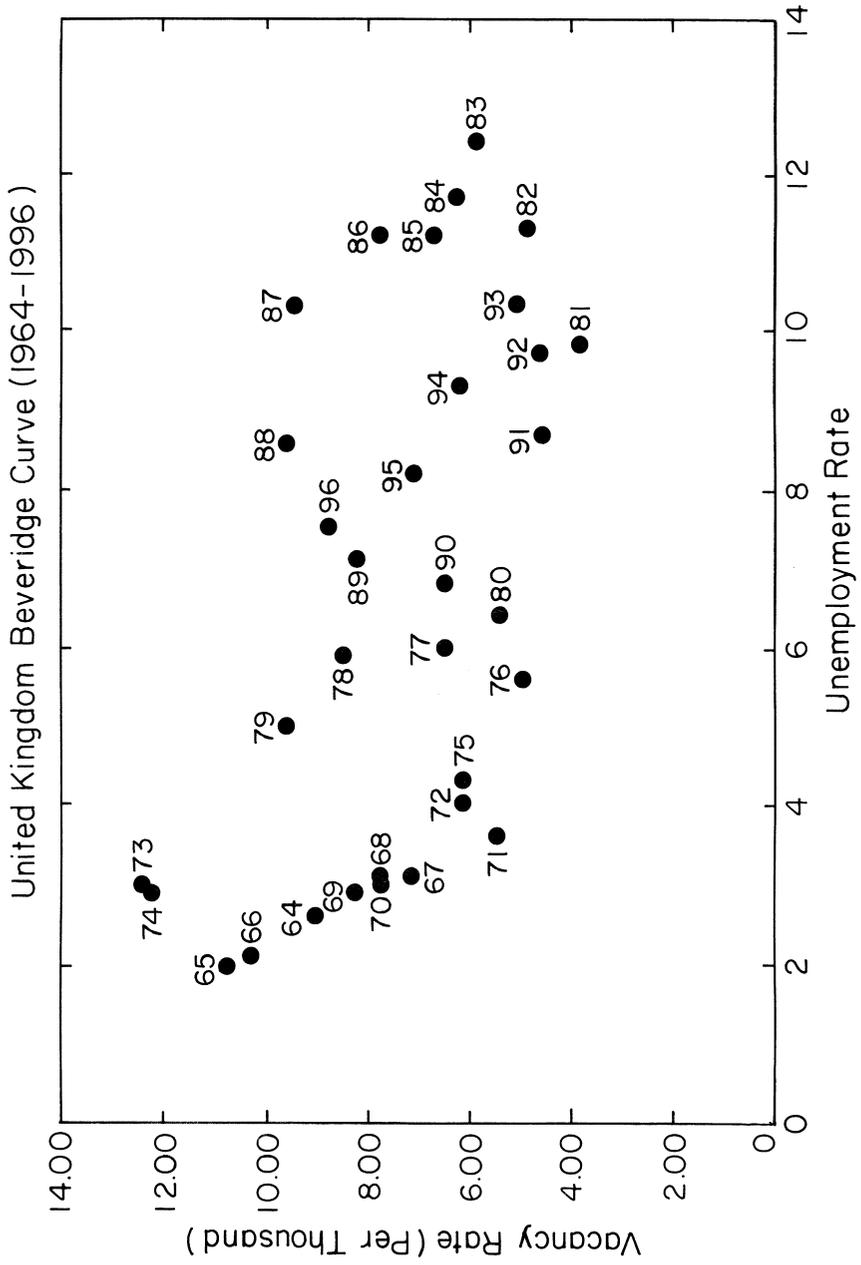


Figure 2.

stable position—at least temporarily—in the 1990s. The initial rightward shift spans almost nine percentage points of unemployment at the extreme; and the reversion to the left takes about four percentage points back. A vacancy rate that would have corresponded to three per cent unemployment in the 1960s is accompanied by roughly eight per cent unemployment in the 1990s. This is obviously a much bigger change than can be inferred in the US.

There is, however, an underlying similarity in timing. In both cases the adverse shift of the Beveridge curve begins around 1972 or 1973. The temptation is strong to identify it in the UK, as in the US, with the slow adjustment of wages to the productivity slowdown that began in those years. (It might once have been thought that the first OPEC oil shock was the source of the maladjustment of wages, or even of the productivity slowdown itself. But that idea has lost whatever plausibility it ever had, if only because the later fall in the real price of oil had no corresponding effect. In any case, the rightward shift of the Beveridge curve seems to have begun a bit too early to be explained in that way.)

The leftward migration of the Beveridge curve also dates from 1987 in both countries. That coincidence might offer a hint as to the underlying cause. But I would prefer to leave that inference to others who know more than I do about the timing of institutional, political, and other changes in the UK and US labour markets that might account for the stories told by the two Beveridge-curve diagrams. Apart from these similarities in timing, there are drastic differences. The two most noticeable are, first, that the adverse shift in the UK was larger and more drawn out in time than the corresponding shift in the US, and, second, that the UK Beveridge curve has reverted only about halfway back to its initial favourable location, whereas in the US the 1960s and the 1990s seem to look alike. It will take a knowledgeable combination of formal analysis and local anecdote to account for those differences. My immediate interest lies elsewhere, and especially in the contrast with the corresponding developments in France and Germany.

For that we can look at the third and fourth Beveridge-curve diagrams (Figs. 3 and 4), which are in fact very much like each other and very different from the preceding ones for the US and the UK. In both France and Germany there is a suggestion of a vertical portion of the curve at the extreme left. This is what one would expect to see if there were a minimal level of frictional unemployment necessary for the labour market to function at all; it would reflect entry and exit from the labour force, turnover from one job to another, and so on. The diagrams make it look as if that minimal unemployment rate were just under three per cent in France, achieved in the late 1960s and early 1970s, and one per cent in Germany, achieved at exactly the same time.

Then the picture gets more interesting. Something may have happened beginning in 1975, in both countries. But the simple configuration of the data

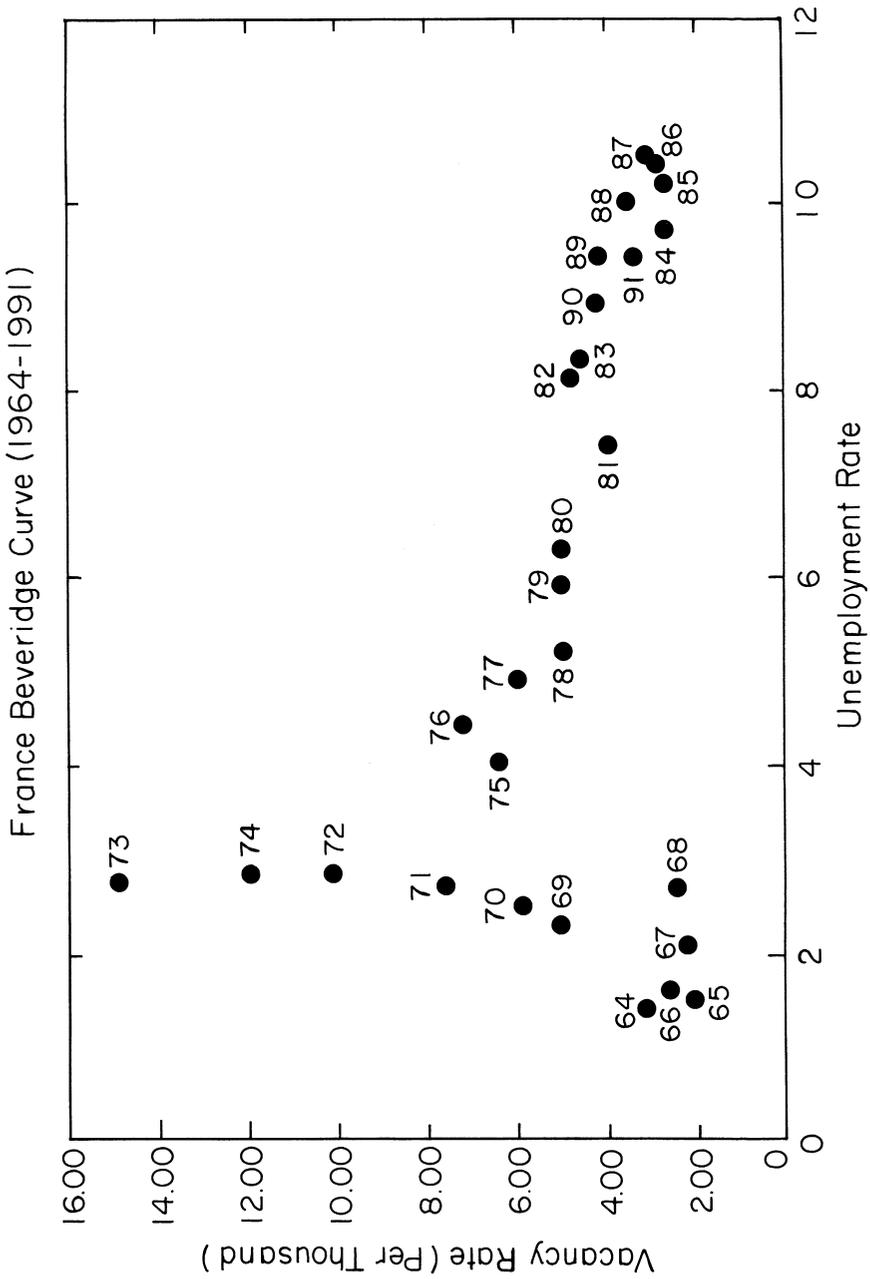


Figure 3.

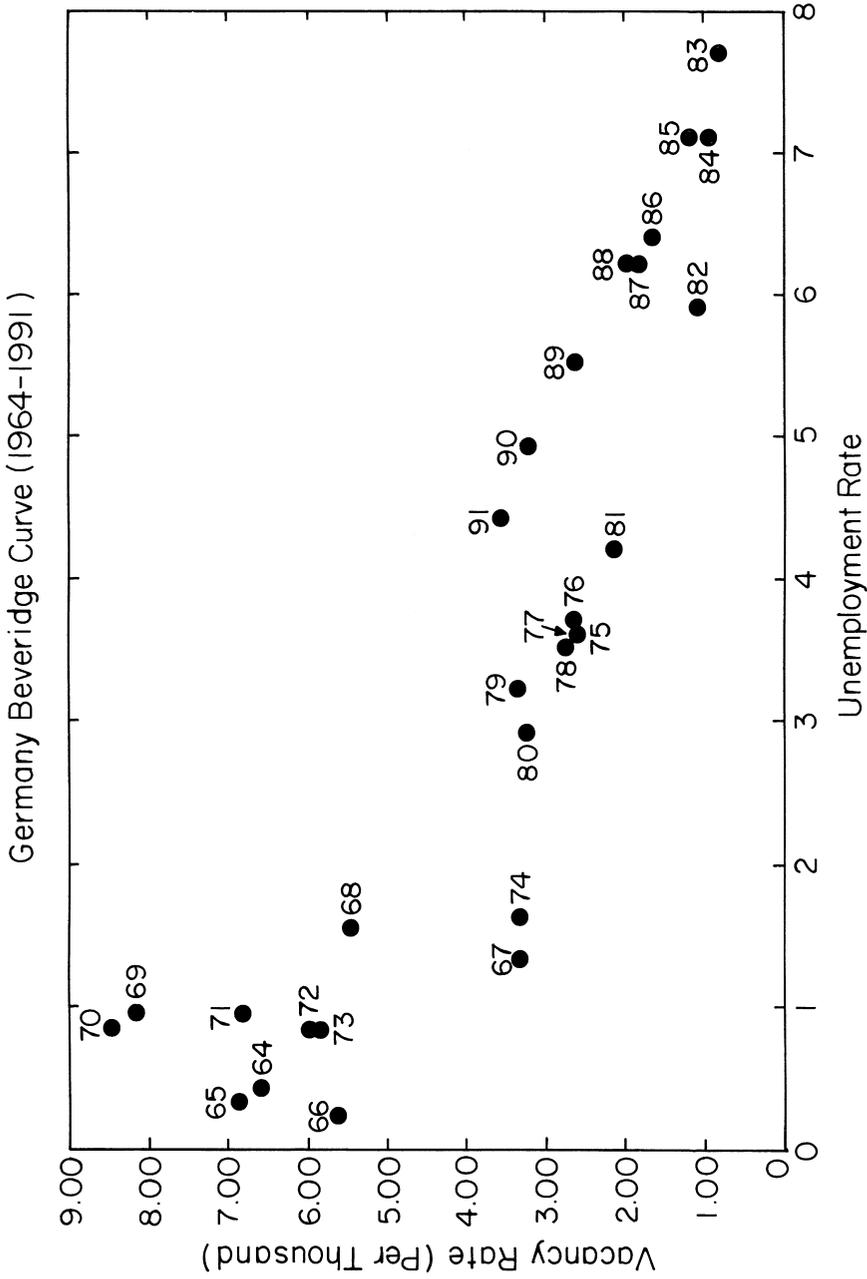


Figure 4.

allows two interpretations. One is that there was a small rightward shift of the Beveridge curve in both countries, amounting to about one percentage point of unemployment in France and fractionally more in Germany. The other is that there was no shift at all, and the whole twenty-eight-year period traces out a single, more or less stable, Beveridge curve. In practice, this is a distinction without a difference, because the adverse shift, if there was one, was so small.

In the case of France, moreover, the years 1964–9 are anomalous. The eye could make a case that a significant shift separates the years before and after 1970. But that would seem to have little to do with the period of endemic high unemployment in the 1980s and 1990s.

Unfortunately the interval described in these graphs ends in 1991, because the later data are for various reasons incomparable with the earlier observations. So we cannot look at the 1990s through this particular lens. However the small reductions in unemployment that took place in France between 1986 and 1990 and in Germany between 1983 and 1991 do seem to be traversing much the same Beveridge curve as was traced out in the opposite direction in France between 1980 and 1986 and in Germany between 1983 and 1991. In saying this I am taking account of the normal presumption that evolving data would trace out counter-clockwise loops around the curve representing stationary equilibrium positions.

I have said that these observations are open to slightly different interpretations. But I also have to claim that the main message transmitted by the Beveridge curves transcends these alternatives. That message goes squarely against the cliché that high and persistent European unemployment is entirely or mainly a matter of ‘labour-market rigidities’. It is precisely in France and Germany, where unemployment has been higher and more persistent, that there is no sign of a big adverse shift in the Beveridge curve. It is precisely in the US and the UK, where unemployment has been at least more variable and, in the case of the US, lower, that one can detect a substantial adverse shift, followed by a favourable one.

To the extent that the location of the Beveridge curve is a reasonable summary for the degree of labour-market rigidity, the large continental economies do not seem to have suffered from noticeably more rigid labour markets during the high-unemployment 1980s than they did in the low-unemployment 1970s. In fact what stands out from the pictures for France and Germany is the depressed level of the vacancy variable. It is a pity not to have comparable data for the last five years.

In the case of Germany, where the data now include the Ostländer, one can at least say that there is no indication of a rebound in vacancies. The case of France is even less clear. There is a new series of ‘new job vacancies’ and it has risen smartly since 1991; but this sounds like a measure of ‘job creation’ and it is impossible to interpret it in isolation from information about job

destruction and pre-existing vacancies. It is a reasonable judgement that the major difference between France and Germany now and in the early 1970s is that the demand for labour is now much weaker. It is not reasonable to blame that large increase in unemployment on *worsened* labour-market rigidity.

I think that the evidence just presented is very strong, but probably not conclusive. That is because the location of the Beveridge curve cannot be a complete summary of the degree of labour-market rigidity at a given time and place. In particular, one of the factors underlying the generation of vacancies is the intensity of job-creation through the appearance of new firms and the expansion or transformation of old ones. It is certainly possible that job creation could be inhibited by apprehensiveness about the working of the labour market. As an example, limitations on their ability to discharge workers if sales expectations are disappointed would surely make employers less eager to create vacancies and hire workers. (Whether this mechanism is quantitatively important is another matter.) In principle this side-effect could be seen as an adverse shift in the Beveridge curve. But one must allow for the possibility that the same effect could be confounded with a movement along the Beveridge curve in the data for France and Germany.

To explore this possibility, I turn to an altogether different kind of evidence. A year or two ago the research branch of the McKinsey consulting firm (it is called the McKinsey Global Institute) conducted an extensive study of economic performance in France and Germany. A handful of academic economists (of whom I was one) participated in that study as an active advisory committee; but the part of the work that I need to emphasise here was done mainly by McKinsey's own people, each contributing experience and expertise in a particular industry. This part of the study consisted of six detailed industrial case studies, each comparing the performance of a particular industry in France and Germany with the same industry in whatever country was the world champion in productivity, usually the US but Japan in one instance and the Netherlands in another. The industries studied were automobiles, housebuilding, telecommunications, retail trade, consumer banking, and computer software.

The notion of 'economic performance' is not self-explanatory. In the McKinsey study a national industry got good grades for high productivity—compared with the benchmark—and for high employment. It is important that in most cases the two go together. The country with the highest productivity in any particular industry is usually the country that has created the most jobs in that industry, and this is true even in industries like telecommunications where it cannot be thought that the benchmark country has drained jobs from others. No economist will be surprised at this finding; it is worth mentioning only because political opinion in Europe often seems to embrace a refined version of the old Luddite fallacy.

These remarks are just by way of introduction. The relevant and important part of the McKinsey case studies is that in each instance the team tried to think explicitly about causality. If a particular industry in France or Germany has been deficient in productivity or in job creation, *why* has it not done better? What factors in the institutional environment and what features of economic policy have caused the shortfall in productivity and in job creation? This part of the analysis was generally not mechanical or even quantitative; it relied on the experience and judgement of the consultants and their informants in each industry, but the advisory committee was able to insist that these judgements be backed up by concrete detail. We would not have tolerated arguments that violated economic logic. This is not the way I am used to arriving at conclusions. But I have to say that I found the process enlightening. The case I am trying to make can only benefit from finding that altogether different approaches lead in a common direction.

In the course of this and earlier studies, the McKinsey group has worked out a standard protocol for characterising causal influences on a given industry in a given place. It is a sort of vertical checklist, beginning with the macroeconomic environment and descending through externally imposed product-market regulations and institutions, through the constraints and costs that stem from the organisation of the labour and capital markets, including the incentives and limitations that arise from pressures on corporate governance, to the intensity of competition and the degree of exposure to industrial best practice, and ending with details of the production process and the organisation of functions and tasks. In the end, each of these potential causal factors is classified as being important, merely secondary, or not significant at all in distinguishing the productivity and job-creation performance of this national industry from the corresponding benchmark. I want to report on these conclusions.

I will describe the conclusions case by case, before summarising the overall message. In the automobile industry, the only mature manufacturing industry among the six, Japan is the benchmark. Comparing France and Germany with Japan, the group finds that restrictive work rules are a factor of only secondary importance, and differences in labour costs are negligible. The important causal factors come from the limited exposure of the European industry to competition, and from inferior management of operations.

In housebuilding, it is found that labour-market factors are essentially insignificant in explaining the productivity shortfall of France and Germany compared with the Netherlands. Differences in product regulations and in internal organisation are far more important. Germany does have higher labour costs than the other countries, including France, and these might be a secondary factor if benchmark productivity were achieved.

In the telecommunications sector, the conclusion is that restrictive work rules are a secondary causal factor in productivity comparisons. The important

causal factors are elsewhere. Independent of productivity differences, the French and German industries generate less output and employment than they might. This surely has nothing to do with the labour market and everything to do with competitive intensity, pricing, and marketing effort.

In consumer banking, the European productivity shortfall has little or nothing to do with the labour market, and much more to do with forces arising in product markets and in internal management. Nor does the labour market have a significant effect on the output of consumer-banking services, conditional on productivity. Differences in job-creation arise elsewhere.

In retailing the productivity differences are small, but the underlying facts are more complicated. An important part of the picture is that high minimum wages in France and Germany induce stores to avoid high-service, high-employment formats. Thus the European industry does not provide the large number of low-wage jobs that it does in the US.

In the software industry, labour-market flexibility is an insignificant factor. All the action in job-creation is in product markets and scale effects.

I have compressed a detailed and sophisticated report by McKinsey into a couple of slogans. But I have to summarise even those. There are a few contexts in which labour-market factors are a significant influence on the number and kind of jobs created. But the bottom line is clearly that these case studies strongly confirm the inadequacy (to put it mildly) of the standard litany that places the blame for low employment in Europe squarely on the inflexibilities of the labour market. It turns out that practised observers of the industrial scene, when they come down to careful, structured evaluation, do not classify labour-market rigidities as an important causal factor in the failure of (at least) these six industries to create more jobs. One might guess—naively, I think—that the observers in question, consultants and business insiders, might normally be disposed to single out the labour market for criticism. The fact that they do not do so lends a little more weight to the conclusion.

This kind of evidence reinforces the interpretation of the Beveridge curves suggested earlier. The likelihood is that France and Germany have moved to high-unemployment regimes by sliding *along* their Beveridge curves, and not as victims of adverse shifts in their Beveridge curves. The implied weakness in job creation is most likely the result of excessive and anti-competitive product-market regulation, restrictive macroeconomic policy, especially monetary policy, and inadequate discipline from the capital markets. This is quite different from the conventional picture.

None of this is to deny that European labour markets are in fact highly regulated and spotted with rigidities. Even if these deviations from a pure spot market are not the main source of the long spell of high unemployment, they may still be a source of real cost to the European economy. Anything that limits the ability of firms to adapt to changed circumstances is a possible

source of inefficiency. So one might prescribe an attempt to achieve greater labour-market flexibility even if this were not expected to have any durable effect on the level of unemployment. There are some qualifications, however.

Flexibility also entails some costs of its own, and they are worth mentioning because they seem to be neglected in current debates, despite their obviousness. A job provides not only a regular wage but also some security of income. It is not far-fetched to simplify by imagining each job to be characterised by its wage and by some measure of its permanence, like its expected duration. Jobs have many other relevant characteristics as well, ranging from safety to sociability, but I would like to focus on security. Employees clearly value both aspects of a job (as who should know better than tenured academics). On the other side of the labour market, both aspects of a job represent costs to employers. The provision of job security is costly precisely because it limits the employer's freedom of action in adapting to exogenous changes in the market environment. One would expect both parties to an employment contract to be willing on some terms to trade off wages against security. One can see this happening in collective bargaining, with further by-play among differently situated workers on that side of the bargain. (Frank Hahn and I have made a first stab at modelling this situation; see pages 95–101 of our *Critical Essay*.)

One obstacle to this line of thought is that it is hard to contract for job security *per se*. The relevant characteristic of a job is its *ex ante* expected duration, not its *ex post* actual duration. It is almost impossible to say whether any concrete act of 'downsizing' represents (a) the occurrence of a contingency whose possibility was foreseen in the original explicit or implicit contract, or (b) an attempt on the part of the firm to shift risk to its employees in violation of the explicit or implicit contract. Whichever is true, the firm will always plead necessity and the workers will resent the *ex post* application of a rule, whether or not it was contemplated *ex ante*.

There are no doubt sophisticated ways to achieve incentive compatibility in this context. In actual fact, however, one imagines that employment contracts try to meet this difficulty by imposing inefficient limitations on the firm's flexibility (for instance by regulating outsourcing). To repeal many or all of such arrangements suddenly by legislation amounts to a shift of risk from a firm to its employees. This is not self-evidently a reasonable thing to do. Workers are almost certainly more risk-averse than the firms that employ them, since they have practically no opportunity for diversification and no possibility to claim compensation through the tax system for a capital loss on job-specific human capital.

The welfare economics of job security is a complicated subject that deserves more analysis than it has had, and surely more than I have tried to provide here. I introduced the subject only to make a point about economic

policy. If pure unadulterated labour-market reform is unlikely to create a substantial increase in employment, then the main reason for doing it is the anticipated gain in productive efficiency, however large that may be. But if we respect the wage earner's desire for job security, and it seems at least as respectable as anyone's desire for fast cars or fat-free desserts, then an improvement in productive efficiency gained that way is not a Pareto-improvement. More labour-market flexibility may still be worth having—and I think it is—but then the losers have a claim in equity to some compensation. The trick is to find a form of compensation that does not cancel the initial gain in labour-market flexibility. Some forms of job-protection, like those described in the insider–outsider literature, provide job security for incumbent workers at the expense of the excumbents. That is not what we are after. But a bit of ingenuity might get somewhere. I call attention to the fact that this is a far cry from your basic European central banker's folk-theorem.

Discussion

Charles Bean, *London School of Economics*

Like Apple Pie and Motherhood, it seems that labour-market 'flexibility' must be a desirable feature of an economy. Yet too often the argument that greater labour-market flexibility is required to tackle Europe's unemployment problem is accepted uncritically. Professor Solow—whom I had the very great honour to study under in my youth—makes some telling and pertinent criticisms of this conventional wisdom. I have much sympathy with his general line of argument that the role of labour-market institutions has been overplayed in the public debate and that of product-market imperfections underplayed. However, I would not want to downplay the role of the labour market quite as much as he does.

Bob first argues that excessively high real wages cannot be the problem in Europe since the wage share has been falling during the last fifteen years or so. Unfortunately this does not imply that factors leading to excessive 'wage push', such as over-generous unemployment benefits, are not the cause of the unemployment problem. In the very short-run it is probably true that excessive wage pressure shows up as an increased wage share. However, as firms cut back on employment the marginal and average product of labour will rise so pushing the wage share back down; in fact with a Cobb-Douglas production technology one would observe no change in the wage share at all. Furthermore the resulting decline in profitability will discourage investment, thus reducing the demand for labour below what it would otherwise be (this is a process to which my co-discussant, Bob Rowthorn, has also drawn

attention). The fall in labour demand, and consequent rise in unemployment, will then lead to a fall in equilibrium wages. Indeed, if there are constant returns to scale in labour and capital together, the equilibrium real wage in the long-run will be determined *entirely* by the state of technical knowledge and the user cost of capital (determined largely in the global capital market); wage-push factors thus show up in higher unemployment, but not in either a higher wage share or a higher real wage. I conclude that one can learn little about the causes of unemployment by looking at the movements in these latter variables alone.

Bob then goes on to argue that the Beveridge Curve linking unemployment and vacancies is a useful indicator, and that a reduction in labour-market flexibility should show up as an outward shift of the Beveridge Curve; this is not obviously the case in France and Germany. Now I agree that it is a useful weapon to add to the diagnostic armoury, but it is only half the story and needs to be used carefully. Figure 5 outlines a simple model of the joint determination of unemployment and vacancies (the exposition follows the work of Christopher Pissarides and his co-authors). The number of successful job matches will be increasing in both the number of firms looking for workers and the number of unemployed workers looking for jobs. Hence the pairs of unemployment and vacancies consistent with a constant unemployment rate can be represented by the downward sloping line, UV; above (below) the UV line unemployment will be falling (rising). This is the Beveridge Curve. However, there is a second line, VS (for Vacancy Supply), that tells us how many vacancies firms open at any given unemployment rate, and thus also where along the Beveridge Curve the economy is located. This VS schedule embodies within it both the factors determining labour demand (technology, firing cost, . . .) and wage determination (unemployment benefits, . . .), and is

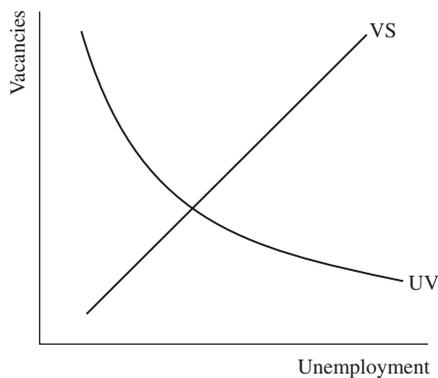


Figure 5.

upward-sloping because high unemployment implies a low level of wage pressure and thus a high propensity of firms to open vacancies.

We can use this apparatus to study the effects of various labour-market institutions and policies. Let me start by noting that increased labour-market 'flexibility' is usually taken to cover a whole raft of policies, including reduced impediments to hiring and especially firing, measures to increase the mobility of labour across both regions and occupations, eliminating restrictions on working time, union bashing, less generous unemployment benefits and lower minimum wages. The effect of each of these is, however, very different and cannot all be collapsed into a simple statement that increased (reduced) flexibility shifts the Beveridge Curve in (out). Thus, an increased mismatch between the skills of the labour force and the needs of firms will indeed just shift UV out as Bob suggests. However, anything that increases wage pressure such as higher minimum wages or aggressive union behaviour will shift VS down. And an increase in the generosity of unemployment benefits would shift both UV out (because it makes the unemployed more choosy about which jobs to accept) and VS down (because it makes the employed more willing to push for higher wages). Consequently the German and French experience is in my view entirely consistent with the roots of the unemployment problem lying within the labour market.

However, Bob supplements his use of the Beveridge Curve with the lessons of the McKinsey comparative study of economic performance. This sort of 'checklist' methodology is somewhat alien to economists, and I have to say that I can understand why, as I find it somewhat difficult to draw out the policy implications. I have already noted that in the long-run real wages are pegged down by the state of technology and the required return on capital, not by labour-market institutions which in general equilibrium only affect unemployment. I thus do not find it altogether surprising that when firms are asked about the most important factor inhibiting job creation that product market and internal management factors might be primary and labour market factors only secondary. However, it might be very difficult for policy to do anything about the primary factors (although I hasten to add that I am all in favour of product market de-regulation in Europe), while the general equilibrium consequences of changes to labour-market institutions may still be quantitatively important. Furthermore even if increased labour-market rigidity is not the *cause* of Europe's higher unemployment, nevertheless labour-market reforms may still be desirable. Put simply, the macroeconomic environment deteriorated in the 70s and 80s because of the productivity slowdown, the oil price shocks and the subsequent disinflation, and Europe's labour-market institutions were simply less well suited to handling this deterioration in the environment than the United States.

Finally let me agree wholeheartedly with a point that Bob makes in relation

to the political feasibility of reform. Most of the rigidities that people say need to be removed or moderated are not accidents of the system, but play a role either in offsetting market failures (mandated severance pay compensates workers who otherwise bear most of the costs of job destruction) or more usually transferring rents (minimum wages may harm the job prospects of the unemployed—although even this is open to debate—but raise the wages of those in work at the cost of shareholders). Reform is difficult because there will always be losers, and in practice it may be difficult to compensate them. Reform will be politically easier to sustain when there are fewer losers, and it will thus be easier to enact reform in a booming economy where jobs are plentiful than in a depressed one. The focus on fiscal consolidation in the run up to monetary union, and the possible future immobilisation of fiscal policy under the Stability Pact, has in my view made it much harder to push through reform than necessary. It would have been far better if the governments of Europe had embarked on labour-market reform *before* setting out on the road to a Single Currency.

Robert Rowthorn

It is always a pleasure to hear a lecture by Robert Solow. He has an unrivalled ability to strip away the extraneous details of any problem and focus on the core issues. This lecture is no exception. Within the space of an hour he has elegantly demolished the conventional wisdom which now dominates European economic thinking. He argues that neither high wages nor rigid labour market practices are the primary cause of the present high unemployment in continental Europe, as exemplified by France and Germany. He concedes that lower wages and more flexible labour markets might help to reduce unemployment, but he regards these as palliatives which do not address the central problem. In his view, the 'weakness in job creation is most likely the result of excess and anti-competitive product-market regulation, restrictive macroeconomic policy, especially monetary policy, and inadequate discipline from capital markets'. With some qualifications, I agree with this unfashionable view, and I am pleased to hear it so well presented by such an eminent economist.

In support of his position Solow presents three pieces of evidence: (1) the relationship between wages, productivity and profits over the past thirty years; (2) the relationship between unemployment and vacancies; and (3) the results of a McKinsey study of industrial performance to which he contributed. Let us take these in reverse order. The McKinsey study found that, in five of the six industries examined, labour-market considerations are not a major factor behind Europe's poor employment record. The one exception is retail distribu-

tion where minimum wage laws inhibit the growth of low-wage jobs on the American pattern. These findings strongly support Solow's views regarding the secondary contribution of working practices to European unemployment. However, their implications for the wages are less clear.

Most of the industries covered by McKinsey pay comparatively high wages even in the USA, and minimum wage laws are of minor relevance to them. It is striking that in the one industry where low-paid employment is important in the USA, the McKinsey study found that minimum wages did inhibit such employment in Europe. If the study had covered a wider range of industries, this finding would probably have been more common. In a cross-section analysis of the OECD countries, Andrew Glyn has identified an inverse relationship between employment and wage dispersion. The employment rate is in general highest in countries with the greatest wage dispersion, in part because very low pay at the bottom end of the earnings spectrum encourages low-productivity employment in labour-intensive sectors. This suggests the following interpretation of modern experience. In most OECD countries the overall demand for labour has been inadequate. In the USA, and to some extent the UK, the response has been to deregulate labour markets, allowing wages to fall and forcing workers to accept whatever job is available. The result has been a proliferation of low-paid jobs mainly in the service sector. In continental Europe this has not been allowed to happen and the wage floor has been maintained, with the consequence that inadequate demand for labour is reflected in overt unemployment. Thus, increased wage dispersion in the USA and higher unemployment in continental Europe are two sides of the same coin. They are the outcome of different responses to a common overall shortage of demand for labour. Opinions may differ as to which response is better.

An important part of Solow's lecture concerns the relationship between vacancies and unemployment, as summarised by the so-called Beveridge curve. He shares the conventional view that the position of this curve is a reflection of labour market rigidity, and that an outward shift in the curve indicates increased rigidity. I agree in principle, but I also think that some caution is required. As Solow himself points out, in the case of the UK there was a huge outward shift in the Beveridge curve between 1975 and 1983, when unemployment increased from four per cent to almost thirteen per cent with no significant change in vacancies. Most of the increase occurred after 1980 when the Thatcher government was in power and busy attacking the unions and dismantling protective legislation. It is difficult to believe that rigidity actually increased during this period, certainly not on a scale sufficient to explain the outward shift in the curve. An alternative explanation is that the economy suffered from pre-existing rigidities which were only exposed when it suffered the shock of a severe crisis. For example, during the period in question, the UK experienced an industrial collapse which destroyed millions of jobs in manu-

facturing, mining, and construction. Traditional industrial areas were severely affected by these losses and the knock on effect on local service activities. The effect was to create a large pool of unemployed workers, who were geographically concentrated, lacking the skills required for other jobs, and reluctant to move. Most of these workers were not able to take advantage of the new opportunities created by the economic boom of the late 1980s, with the result that unemployment remained fairly high despite a pronounced shortage of labour in some areas. This outward shift in the Beveridge curve cannot be ascribed to increased labour market rigidity, but to pre-existing rigidities which inhibited labour mobility and the retraining of workers following a major, regionally and industrially specific, shock. The same is probably true on a less dramatic scale in the USA. It is interesting that large shocks during this period do not seem to have shifted the Beveridge curve in France and Germany. It may be that labour mobility or retraining possibilities were greater in those countries, or the shocks less regionally specific, or regional job creation programmes more effective. Whatever the reason, Solow is right to point out the difference, and to insist on the fact that France and Germany have been sliding along their Beveridge curves, with the implication that what is really at fault is an overall lack of demand for labour in these countries.

This brings me to my final point which concerns wages and profitability. I think that Solow understates the significance of wages for employment. The conventional theory of the Beveridge curve tells us that an increase in real wages may simultaneously reduce the number of vacancies and the number of people in employment, thereby causing the economy to slide down the Beveridge curve. The observed behaviour of unemployment and vacancies in France and Germany is therefore theoretically consistent with the fact that wages are too high. However, this is probably too static an interpretation. The dynamic issue is the effect of wages on profitability, and thereby on investment and future employment. It is here that the wage gap literature, of which Solow is so critical, is relevant. Profits were squeezed during the 1970s because of worker militancy and the failure of wages to absorb the full cost of higher oil and commodity prices. Firms increased prices to protect their profit margins, but this led to accelerating inflation in many countries, and governments eventually responded by raising interest rates and provoking an economic crisis. Real profits fell still further as capacity utilisation declined, and the result was a collapse in investment. In the ensuing recovery, profits were partially restored and investment picked up, but the recovery was not sufficiently strong or prolonged to compensate for the previous shortfall in investment. The situation has been made worse by the restrictive policies applied during the run up to EMU. As a result of these events, most Continental economies are now much too small to provide employment at reasonable pay for those desiring work.

I have not the time to analyse in detail why European investment has been so low, and I shall content myself with the following observations. For much of the past quarter century European economies have been characterised by excess capacity as a result of government anti-inflationary policies, and many of them experienced a decade of low profitability in the 1970s and 1980s. There has been a widespread recovery in profits, but they remain quite low in Germany, especially when compared to the opportunity cost of capital. For example, the net business profit rate in Germany averaged 13.5 per cent over the period 1969–73 as compared to a long-term real interest rate on government bonds of 2.0 per cent. For the period 1989–93 the corresponding figures were 10.8 per cent and 3.7 per cent. Thus the gap between net profits and interest fell from 11.5 per cent to 7.1 per cent. I do not have the equivalent figures for other Continental countries, but the figures on profit shares in manufacturing suggest that net profits have recovered to their pre-Oil Shock levels in most of them. However, long-term real interest rates have been very high for many years, and much of the time firms have been operating with excess capacity, both of which have served to depress investment and prevent countries from making good their long-standing shortage of capital stock. To produce a substantial reduction in continental unemployment requires a period of above average investment and growth. To achieve such a combination requires both expansionary demand policies and a prolonged period of above average profits.

The above observations suggest that wages may still be a factor behind unemployment in Europe through their effect on profits and investment. If wages were lower, profits would be higher and there might be more investment, and ultimately more jobs. Moreover, the *fear* of wage inflation causes governments to adopt restrictive policies that create unemployment and inhibit investment. It may be that such a fear is mistaken, because in the medium term more investment might reduce inflation by increasing productivity and thereby offsetting the inflationary effect of higher wages. If this is the case, governments should be willing to accept a transitory period of higher inflation before new capital stock comes on stream. Of course, to advocate such an approach is hopeless at the moment, because policy makers are still dominated by an inflation phobia which makes them unwilling to risk even a temporary acceleration in inflation as the price of more investment and more jobs. However, if high unemployment continues, political pressures may eventually force the authorities to rethink and start to experiment with more expansionary policies.

To sum up. I agree with the basic thrust of Robert Solow's argument that labour market rigidities are not the primary cause of high unemployment in Europe, but I would give more of a role to wages and the wage-profits nexus than he does. However, this is a minor qualification and I strongly welcome his lucid and stimulating contribution.

