

An Introduction to the Wage Curve

June 1994

David G. Blanchflower
Dartmouth College and NBER

Andrew J. Oswald
London School of Economics

This paper draws partly upon work to be published, in the fall of 1994 by MIT Press, as a research monograph entitled The Wage Curve. The research was funded by the ESRC.

Abstract

This paper describes a statistical regularity or empirical 'law' of economics. Research done in the last few years has uncovered what seems to be a negatively-sloped curve linking the wages earned by employees to the unemployment rate in their region (or industry). The curve emerges from microeconomic earnings equations estimated on international data sets that in some cases use hundreds of thousands or even millions of observations. It holds after controlling for workers' personal and demographic characteristics, for regional characteristics, and for macroeconomic and other aggregate influences. The curvature of the function -- an index of wage flexibility that might be termed the unemployment elasticity of pay -- is apparently approximately the same in each of the fifteen countries in which the curve has been found. That elasticity is -0.1.

An Introduction to the Wage Curve

David G. Blanchflower and Andrew J. Oswald

The market for labor is complex. At first glance it seems to be little more than a swirl of idiosyncratic detail. People's wages are outwardly the product of quirk, precedent, bureaucratic convenience, luck, and historical happenstance. Institutions differ across regions and nations; shocks take diverse forms over time; notions of demand and supply seem here to come confusingly bound up with people's inner notions of fairness. It might be thought that this complexity and randomness would make hopeless the search for logical structure in these kinds of labor-market data. Yet economists have somehow to uncover patterns that lie behind the detail.

This paper describes an empirical discovery -- made with the help of international microeconomic studies covering more than a dozen nations -- about the workings of the labor market. It features two of the variables that most interest policy-makers, namely, the level of unemployment and the level of wages. These are linked, it is argued here, by a curve or formula described by the diagram below.

Figure 1 here

Because the wage curve is not part of conventional economics, most of the chapters in our book, The Wage Curve, are designed as collections of new evidence for it. Data on approximately three and a half million people, from twelve countries, are used in that book's analysis.

There is corroborative evidence for the curve's existence. Some of this microeconomic work covers other developed (and even developing) nations. Although much remains to be understood, the combined picture is, by the standards of modern economics, a consistent one. It is summarized in the pages to come¹.

The relationship depicted in Figure 1 raises natural questions and doubts. What does this curve mean? Is it not unrealistic to believe in laws of behavior in a subject, like economics, that studies human beings? How, in a field where experiments are almost impossible, could there be any convincing scientific basis for this wage curve? Why does the Figure look like a Phillips

Curve with the vertical axis mis-labelled? Should it not, for Harris and Todaro (1970) reasons, slope up and not down? How does the relationship fit into existing knowledge? What causes the curve?

Methods and Microeconomic Data

Any construction rests ultimately upon its methods and raw materials. The modus operandi in research on the wage curve has been to draw upon random samples of workers and establishments. These data come from various nations. When the same statistical methods are used on each nation's sample, the international data sets reveal approximately the same answer, that is, they seem to point again and again to a curve like that in Figure 1.

To be convincing, such an approach needs data of persuasive quality. In natural science, data would normally be collected as part of controlled experiments, and a proposition would not normally be thought of as proved until many different investigations had replicated the essential finding. Yet some science, especially medical science, proceeds in a less orderly way. A good example is the discovery, usually attributed to the Oxford researcher Sir Richard Doll, that smoking raises the probability of contracting lung cancer and other diseases. This fact first emerged from an examination of survey information on British individuals. The data contained an elementary 'natural experiment': some people smoked and some did not. With the help of simple statistical methods, it could be shown that, other things constant, there seemed to be a relationship between smoking and ill-health. Since then the same pattern has, of course, emerged in many other studies, and most Western governments have brought in policies to discourage smoking. Doll's discovery provides an illustration of how non-experimental methods can uncover relationships. Although the topic covered in this paper is far from those tackled in medical science, there are similarities in the underlying methods.

Research on the wage curve cuts across the three fields of labor economics, macroeconomics, and regional economics. In a sense, regions are treated as mini-economies. This provides many more units of observation than in conventional macroeconomics. For the United States, for example, some results focus upon fifty states through the 1980s rather than on

ten aggregate time-series data points. Issues of spatial economics are then impossible to avoid. How does the fact that people can migrate from New Mexico to New Hampshire, or from Aberdeen to Aberystwyth, fit into the picture? What shapes the cross-section structure of the main labor market indicators? These topics are discussed below. Occasionally the analysis looks at industries rather than regions, but the underlying idea, that of uncovering the links between unemployment and earnings by using disaggregated information, remains the same.

Economists in the 1990s are among the first researchers to have at their disposal microeconomic data of a genuinely internationally comparable kind. One such source is the International Social Survey Programme. This is a multi-country collaborative project, described in sources such as Jowell et al (1989), in which survey teams from a dozen nations asked the same questions of random samples of individuals from each of the countries. Assuming that difficulties of common translation can be overcome -- perhaps not an insurmountable obstacle when asking respondents about objectively measurable items like income or numbers of years of education -- a data set of such a kind endows social scientists with a statistical resource of value. Comparable data across nations offer a natural test-bed for scrutinizing theories. This is because deep explanations of economic and social behavior are presumably those that are common to peoples with different histories and institutions. A focus on internationally equivalent microeconomic data sets may mark this research as unusual in the 1990s, but it is not impossible that in the future it will come to seem second-nature.

Uncovering the Curve

In the The Wage Curve, the size of sample varies dramatically from one part of the world to another. The years sampled also vary. For the United States, for example, the analysis draws upon the Current Population Surveys from 1964 to 1991. This provides a sample of more than one and a half million American workers. Some recorded interviews in the 1960s; others did so a quarter of a century later. A similarly large sample is available for South Korea. That country's Occupational Wage Surveys of 1971, 1983 and 1986 offer the investigator information on approximately 1.4 million employees. At the other end of the spectrum, there are results for small

countries like Switzerland and Norway. These data are from the International Social Survey Programme, and come from the late 1980s. In this case, the samples are of less than three thousand people in each nation.

For each person, from each country, the data sets record the individual's level of pay, and his or her personal characteristics. In most instances, the name of the region in which the worker lives is also recorded in the data set. Rates of unemployment vary, as is well known, from one region to another, and from one country to another. The analysis matches those unemployment rates to the sampled individuals. In other words, if Miss X of San Francisco is known to be earning \$12,000 per year in 1981, the unemployment rate in California in that year can be merged into the data set and treated as a variable that is potentially relevant in explaining the level of Miss X's pay. For Britain, similarly, the sample might include information on Mr Y working in Edinburgh in the year 1974. The 1974 unemployment rate for Scotland, therefore, can be added to the data set as a possible influence upon Mr Y's remuneration in that year. Repeating this process many times gives information on a large number of workers' pay levels and, in each case, an associated unemployment rate in the person's region. Some people live in areas of high unemployment; others are surrounded by little joblessness. This variation offers useful statistical information.

What emerges from the data is a pattern linking pay and unemployment. Just as in the case of smoking and lung cancer, it is best to control for many other factors (such as age and gender) when performing the calculations. Once this is done, a locus of points is visible. This takes the form of a downward-sloping convex curve in wage/unemployment space. A worker who is employed in an area of high unemployment earns less than an identical individual who works in a region with low joblessness.

Perhaps surprisingly, the nature of the relationship appears to be the same in different countries. The wage curve in the United States is very similar to the wage curves in, for example, Britain, Canada and Norway. As a crude characterization of the data, the wage curve is described by the formula

$$\ln w = -0.1 \ln U + \text{other terms}$$

where $\ln w$ is the logarithm of the wage, $\ln U$ is the logarithm of the unemployment rate in the worker's area, and the other terms in the equation are variables for further characteristics of the worker and his or her sector. The equation, which seems to hold in each country, implies that the unemployment elasticity of pay is -0.1. A hypothetical doubling of the local unemployment rate is then associated, consistent with early evidence reviewed in Oswald (1986), with a drop in pay of ten per cent (that is, a fall of one tenth)².

There is another way to think about this finding. When treated in the way just described, the data provide a method for the calculation of an index of wage rigidity. The concept of wage "stickiness" has long been central to much of macroeconomics. It has proved harder, however, to agree on a measure of wage flexibility or inflexibility. The analysis can be seen as offering such a measure. The responsiveness of workers' remuneration to the state of the labor market is captured here by the coefficient on log unemployment in an equation for log earnings. Later results suggest that it is approximately the same in each of the countries considered. This uniformity was not expected by any researcher, and seems remarkable. Despite countries' institutional differences, they exhibit more uniformity in wage flexibility than has been thought. Put more technically, the local unemployment elasticity of pay appears to be roughly the same across nations.

A glance at the relationship plotted in Figure 1 conjures up thoughts of the curve noticed by A.W. Phillips (1958) using nine decades of British aggregate data. The Phillips Curve, however, is not what is being estimated in research on the wage curve. Phillips's construction -- one now in every macro textbook -- is purportedly a locus linking the rate of change of wages to the level of unemployment. The present focus is different.

- (i) The Phillips Curve was proposed as a disequilibrium adjustment mechanism. The wage curve is instead to be thought of as an equilibrium locus that is not, in a useful sense, a description of inherently temporary phenomena or of transitory dynamics.
- (ii) The Phillips Curve links the rate of change of pay to the aggregate unemployment rate. The wage curve links the level of pay to the local unemployment rate.

(iii) The Phillips Curve is traditionally estimated on time-series macroeconomic data. The wage curve is estimated on pooled microeconomic data.

It would be misleading, nevertheless, to view later results as divorced from the tradition that Phillips began. Phillips's underlying concern, to understand the macroeconomic influence of joblessness upon wage-setting, is continued here. Moreover, modern error-correction models are an attempt to distinguish between dynamics and the properties of long-run equilibrium. They suggest that it is valuable to specify Phillips-style equations in a way -- namely with an extra wage-level term on the right -- in which the steady-state can exhibit a connection between wage levels and unemployment levels. Sargan (1964) may have been the first such paper, although, because of its poorly defined standard errors, the substantive results about wage determination are not compelling.

In the new literature, lagged dependent variables often enter microeconometric earnings equations with a coefficient close to zero. This matters, because, contrary to what almost every undergraduate is taught, it is then best to think of unemployment as explaining pay rather than pay inflation. Later results are consistent with the belief that it is poor data and the omission of suitable control variables -- including those for fixed effects -- that leads to spuriously large coefficients on lagged dependent wage variables. The idea of a Phillips Curve may be inherently wrong³. It may be a mirage produced by aggregation error in macroeconomic data. If this disturbing vision proves correct, macroeconomists will probably have to turn to microeconomic data.

A Summary of Key Findings

The main conclusions of research on the wage curve are the following.

1. *A wage curve exists. Employees who work in areas of high unemployment earn less, other things constant, than those who are surrounded by low unemployment.*
2. *A picture of this downward-sloping locus can be drawn with either regional unemployment or industry unemployment on the horizontal axis.*
3. *Approximately the same curve holds in data from the USA, Great Britain, Germany, Canada, Austria, Holland, Switzerland, Korea, Norway, Ireland, Italy, Japan, Australia, the Cote d'Ivoire,*

Sweden and India.

4. *The idea of a Phillips Curve may be inherently wrong. Using microeconomic data, and controlling for fixed effects, the autoregression found in macroeconomic wage equations tends to disappear.*
5. *In the countries studied in this book, the estimated unemployment elasticity of pay is approximately -0.1. This uniformity runs counter to orthodox teaching (based on time-series analysis).*
6. *Harris-Todaro orthodoxy, which says that regional wages are positively correlated with regional unemployment, is misleading.*
7. *Bargaining and efficiency wage models, because they give the correct prediction, are consistent with the pattern discovered in the data.*
8. *It is difficult to see how the wage curve can be compatible with the textbook competitive model of the labor market.*
9. *The wage curve helps to provide the missing empirical foundation -- a flat quasi-supply function -- for a new class of non-competitive macroeconomic models. The evidence suggests that it is not a labor supply curve.*
10. *Microeconomic data drawn from internationally comparable random samples provide economists with a rich resource for testing hypotheses.*

In this research, the local rate of unemployment is usually entered as a logarithm in wage equations, so that, because the dependent variable is also in logs, its coefficient can be read off as an elasticity. A representative table of unemployment elasticities derived from The Wage Curve is given in Table 1. They measure responsiveness of pay in different nations.

Britain, Canada and the USA produce similar results to one another. The estimates of the unemployment elasticity of pay lie in a range from -0.08 to -0.11 in almost all specifications. For US data it is especially important to control for regional fixed effects, that is, to include a set of region dummies or to difference the data. Once this is done, Hall's (1970,1972) predicted positive spatial correlation between pay and joblessness -- the result in theory of compensating differentials

across regions -- becomes strongly negative. This raises doubts about conventional wisdom in US labor and regional economics.

The twelve coefficients summarized in Table 1 are negative and in most cases well-defined. For some nations there are few years of data across which to pool, and it is then to be expected that the inclusion of regional dummies will lead to low t-statistics. Southern Ireland is the analysis's only real outlier. Its coefficient is so unstable that the results were almost left out; they should be treated sceptically. Korea, for which an industry wage curve alone can be estimated, has a low unemployment elasticity of -0.04. Future research, no doubt, will aim to chart divergences from the minus-point-one rule.

Some economists have estimated wage curves for other nations. Montgomery (1993) obtains an unemployment elasticity of pay for Japan of -0.15 or less. He finds little sign -- contrary to popular preconceptions -- that Japan's labor market is unusually flexible. Rebick's (1990) Japanese results are similar. Hoddinott (1993) produces an estimate of -0.13 for the African country the Cote d'Ivoire. This is doubly interesting, because it is one of the first estimates of the local unemployment elasticity of pay for a developing country, and because the calculated number is so like that in the industrialized West. Bhalotra (1993) finds a negative effect for India. Although the coefficient is not always well-determined, once regional dummies are included the estimated elasticity at the mean is, somewhat remarkably, again -0.13. Edin and Holmlund (1989), Holmlund and Skedinger (1990) and Edin, Holmlund and Ostros (1993) estimate a small but well-determined negative coefficient for regional unemployment in Sweden. Further results are in Groot et al (1992) for the Netherlands, Wagner (1994) for Germany, and Winter-Ebmer (1994) for Austria.

To economists raised on the Harris-Todaro (1970) and Hall (1970, 1972) models, the results given here may be a surprise. Those models predict that, by the requirement for a spatial equilibrium in which all regions provide the same expected utility, high-unemployment areas ought also to be high-wage regions. In other words, to ensure that nasty and nice features of regions cancel out, the wage locus should slope up. This hypothesis is as decisively rejected by the

international microeconomic data as it is possible to imagine. That does not mean, however, that the idea of compensating differentials is wrong or irrelevant. Movements in actual wages can be negatively correlated with movements in actual unemployment while at the same time 'permanent' unemployment, put loosely, is positively related to 'permanent' wages. Migration is presumably driven by these permanent values. The mistake that some early empirical literature made was to omit controls for regional fixed effects (that is, to leave regional dummies out of the estimated wage equations).

Possible Explanations

The evidence suggests that the wage curve is not the product of some mechanical error, or of a compositional omission or bias, or of some failure to control for hours of work or the automatic drops in income of those who lose jobs. It does not seem to be a special result that can be discerned only in particular time periods or places. Therefore, it provides a theoretical challenge.

Most economists react to an empirical phenomenon by reaching initially for supply-and-demand analysis. Could a framework of that kind explain the wage curve of Figure 1? At first glance it could not. Unemployment in the neoclassical model of the labor market is given by the difference between supply and demand. If the wage exceeds the market-clearing level, the numbers of individuals who wish to supply their labor exceeds the number whom firms wish to employ. Unemployment results. The higher the wage, the greater the degree of unemployment. In a framework of the textbook kind, where unemployment is the gap between a supply curve of labor and a demand curve for labor, wages and unemployment are positively associated. This is the reverse of the curve in Figure 1.

The demand-and-supply model of the labor market cannot be dismissed quite this easily. Unemployment might be viewed as a voluntary choice by workers. Many commentators have pointed out that this is difficult to square with the apparent distress shown by people when they lose their jobs. Nevertheless, there is a long tradition within economics that questions the usefulness of the concept of involuntary unemployment. Thus a possible way to make Figure 1

consistent with a demand-and-supply analysis is to suggest that the curve in the Figure is a kind of labor supply function. The key assumption then is, of necessity, that unemployment is in some sense the negative of employment. If the potential labor force is fixed at some number greater than the level of employment on a labor supply curve, unemployment does not have to be thought of as a gap between supply and demand. It could be viewed as the gap between labor supply and a fixed labor force.

Unfortunately, logical difficulties then spring up. In particular, the unemployed, according to this definition, are not part of the labor supply curve. They are not offering their labor for sale. This seems inconsistent with the technical definition of unemployment used by most government survey agencies, and perhaps runs counter to common sense. However, neoclassical analysts are still likely to see Figure 1 as a mis-measured version of the upward-sloping supply curve of labor found in most textbooks.

The possibility of a labor-supply interpretation of the wage curve of Figure 1 is considered in detail in some of the research on the wage curve. This is done by comparing the explanatory power, within a wage equation, of both unemployment and conventional measures of the supply of labor. If unemployment is a mis-measured version of a normal labor supply relationship, it should perform statistically less well than supply variables like the participation rate or the employment/population rate. Empirically, however, it does not. The evidence given in The Wage Curve does not offer support for the idea that the negative correlation between pay and unemployment is explained by a labor supply function. The book argues that the demand-and-supply framework is the wrong way to think about the labor market. As Robert Solow's 1989 lectures at Berkeley suggest, there may be something special about labor as a commodity, and therefore about the labor market itself (p.3, Solow 1990).

If the competitive-market model gives an unconvincing rationalization of the wage curve, non-competitive accounts of the labor market have to be considered. In a bargaining model, a high degree of joblessness in the surrounding labor market might be expected to reduce the ability of workers to claim a large share of the surplus to be divided. Non-economists probably think such

an idea too obvious to be worth pondering, but economists are required to outline the mechanism at work.

A possible story -- one theory of the wage curve -- is that outside unemployment frightens workers. This may be because if something goes wrong, and bargaining reaches permanent impasse, the individual workers involved will need to obtain other jobs. Finding jobs is likely to be harder when the local labor market is depressed. Therefore, although some details of the process here remain cloudy, rising joblessness might be thought to spawn declining levels of pay. A variant on this species of explanation relies on the explicit assumption of a trade union that worries about both its employed members and its unemployed members. High unemployment means that more of its members are likely to be without work, and that an employed member who is dismissed or laid-off will have difficulty finding new employment. An increase in unemployment may then tilt the union's preferences towards a greater concern with the number of jobs. If this implies a reduced concern for pay, or at least a slightly smaller weight on the target of high remuneration for union members, a lower negotiated level of pay could result.

While a bargaining or union approach might thereby render intelligible the pattern in the data in Figure 1, it could be viewed as inappropriate in many settings outside Western Europe and even in some industrial sectors within that part of the globe. Unionism is not pandemic. In the United States, in particular, the great majority of those who hold jobs do not belong to a trade union. This does not make bargaining theory irrelevant, but it raises some doubts about its pertinence.

Another way to provide an intellectual rationale for the wage curve is by appealing to efficiency wage theory. This approach is intrinsically non-union, so it is potentially suitable for economies that are more like the US's than Sweden's. The well-known characteristic of efficiency wage analysis is that firms set pay in an environment where the wage influences productivity. Shapiro and Stiglitz (1984) is an archetypal case. In equilibrium, firms try to maximize profits, and workers choose how hard to work. The outside rate of unemployment plays a central role, because it determines the ease with which a sacked worker can get another job. In a highly

depressed labor market, employees are frightened of losing their jobs, and so put in high effort even if pay is comparatively low. Put differently, a marginal rise in unemployment leads to a corresponding marginal fall in the level of wages. The reason is that firms can reduce pay slightly while still maintaining a motivated workforce. Unemployment is a discipline device: when it is high the generosity of workers' remuneration will, via a no-shirking condition, be low. Hence there is an efficiency wage interpretation of Figure 1's pattern.

A further possibility exists. It is a half-way house between a union model and competitive theory. A "labor contract" model might be able to explain the negative correlation between wages and unemployment. In some versions of labor contract theory, wages and employment are positively correlated. A contract curve slopes upwards in wage/employment space because when the wage rate is high it is sensible for the workers and their firm to ensure that as many as possible of those in the labor pool have jobs. To those economists raised on neoclassical labor demand theory, this argument may have a strange ring to it. Should not employment be low when the price of labor is high? In this contractual framework, the answer is that it should not. An efficient labor contract maximizes the joint welfare of employer and employees. Unemployed workers contribute only a little to the surplus earned by the combination. The higher the level of pay, the more desirable it is, from the combination's point of view, to have extra individuals in work and fewer individuals drawing unemployment benefit. An upward-sloping contract curve here acts as a quasi-supply curve. It is therefore worth considering as another potential source of a negatively-inclined relationship between pay and unemployment. There is a variant on this. Consider a situation in which demand shocks occur randomly and the firm has to design a remuneration package that will both make it money and attract sufficient employees. If the firm dislikes risk, it will wish the wage to rise in good times and to fall in bad times. When there is a boom, many workers are employed. In a slump, some are laid off. There is the basis here for a model in which pay and unemployment are negatively related.

The wage curve in labor-contract theories is not truly a causal relationship. The variables' apparent interdependence is a trick of the eye. This approach therefore suffers one of the

disadvantages of the competitive demand-and-supply apparatus. It is inherently about employment rather than unemployment. To be convincing, some additional step is required. During the last few years, theoretical macroeconomics has begun to change. A generation of models has sprung up in which an aggregate wage curve (in this paper's terminology) is the primary distinguishing feature. The exact history of this research current is discussed in Layard, Nickell and Jackman (1991) and Phelps (1990). Influential contributions include Rowthorn (1977), Shapiro and Stiglitz (1984), and Layard and Nickell (1986). The same kind of model was presented by David Soskice in lectures at Oxford University in the 1970s. Rowthorn (1977) talks of the level of demand as a "discipline" that leads to consistency between workers' desire for higher wages and firms' desire for lower wages, and of equilibrium as requiring "the creation of a permanent reserve of unemployed labour ... higher than the frictional amount envisaged by the monetarists " (p.237). Layard and Nickell (1986) neatly term this process the battle of the markups. A textbook introduction to these ideas is contained in Carlin and Soskice (1990). The gist of all this work is that, in the 1980s, "a surrogate employment supply curve, or equilibrium wage curve, was born..." (p.1004 Phelps, 1992). A recent book by Lindbeck (1993) adopts the same approach and uses the assumption that there is a wage-setting curve, different from and to the left of conventional atomistic supply, that slopes upwards in real-wage/employment space. Another model in which a central part is played by the same form of wage equation is the fairness approach described in Akerlof and Yellen (1990).

As Shapiro and Stiglitz (1984) and Layard and Nickell (1986) make clear, the novel aspect of these models is not their assumptions about labor demand, which are the standard ones. Rather, it is that the models replace the conventional labor supply curve with a wage-fixing function. This allows the theories to explain, or at least to be consistent with, both involuntary unemployment and the paradoxical fact -- noted by Greenwald and Stiglitz (1993) and others -- that real wages fluctuate little over the cycle while the long run supply of workers appears to be close to vertical. The crucial constituent part in this way of thinking is a negatively-sloped curve linking the level of pay to the level of unemployment. As this paper tries to make clear, there is now

microeconomic evidence for such a curve.

Criticisms and Interpretation

Three potential criticisms are taken up in The Wage Curve book. First, one chapter checks the hypothesis that, because unemployment depends upon pay (perhaps through the functioning of a downward-sloping labor demand curve), OLS estimation of the wage curve is flawed by a lack of allowance for simultaneity bias. Attractive though such thinking is theoretically, little support for it can be found empirically. In US data, instrumenting regional unemployment by weather variables, military spending, industry mix, or lagged unemployment produces estimates of the unemployment elasticity of pay that are only fractionally higher than those from OLS methods. Unemployment apparently has the characteristics of a predetermined variable. Second, especially for Britain and the USA, it is necessary to correct the standard errors for bias caused, in estimation where the independent variable is more highly aggregated than the dependent variable, by common group effects. The nature of the difficulty, which is routinely ignored in empirical research, is explained in sources such as Moulton (1986). When a correction is done, the substantive findings remain unchanged. Third, for only one country are regional consumer price data available, but theory suggests that it is real wages that are depressed by local unemployment. Although undesirable, there are reasons to think that this problem is not too serious. Controlling for regional prices in the British case leaves the wage curve intact. Nominal wages are likely to be sufficient whenever year dummies and regional dummies can be, as for most countries, included in the regression equations. The evidence for an industry wage curve, moreover, is presumably immune to the criticism that there are not geographical deflators. Finally, if prices depend dominantly upon wages and a constant, they can be substituted out to leave a nominal reduced-form wage equation.

A few economists, particularly those convinced that it makes sense to use the same model when studying the labor market as when studying the market for strawberries, may look with surprise at the correlation reported in the paper. These readers will see the wage curve as an inexplicable reduced-form equation. Where, they are likely to ask, are the structural neoclassical model and the supply-demand specification that any satisfactory research must, surely, contain?

This is a sensible reaction, and one that has probably helped the discipline of economics to advance. But in the 1990s it is dictated by habit. It carries with it the danger that the idea of the wage curve will be pushed aside because it is outside our pre-set mental pictures. First, the pattern documented in this paper is not going to go away if we close our eyes. Second, the regularity in the international data seems too uniform to be plausibly interpreted as a chance correlation, some statistical serendipity, caused by an identification or measurement problem. Third, the research's purpose is, among other things, to question the notion that labor market outcomes are usefully viewed as cutting-points on graphs inscribed with demand functions and supply functions. It does not make sense, therefore, to glance into the research and ask to be shown the underlying demand-and-supply framework behind the data's configuration. The wage curve replaces the labor supply curve.

Policy-makers are also likely to ask what there is to learn from a wage curve. That is a legitimate question, but not yet one that this research can answer in a detailed way. The analysis is mildly consistent with non-competitive models that deliver the view that governments should see involuntary unemployment as existing in equilibrium, and that they should consider whether it is possible to alter beneficially the slope of their nation's wage curve. At this juncture, nothing stronger can be said.

Conclusions

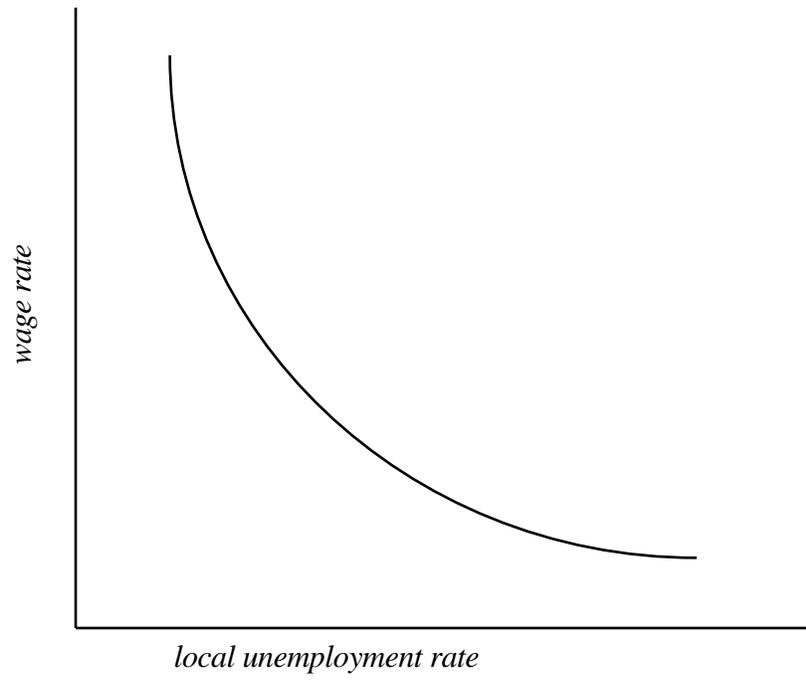
Economics, it might be argued, is either about the discovery of empirical laws or about nothing at all. In 1994, this view is probably not fashionable. Most economists have been raised to think that empirical laws are for chemists and biologists but not for them. Some economists may be unaware that there are now randomly-sampled, multi-country data sets that offer remarkable empirical opportunities to those willing to behave like inductive scientists.

This paper documents the existence of what is apparently an empirical regularity, or statistical law, of economics. Random samples of individuals, from many nations, are used. These samples record on magnetic tape a large amount of personal statistical information about (literally) millions of people, and in some cases plants, across the world. Economics is only

starting to explore the statistical pictures that these can paint. In this kind of research, the methodology is not especially different from that used in much of, for example, medical science.

There seems to be a curve linking wages to the local rate of unemployment. It slopes down: higher joblessness in a region or industry means lower pay in that region or industry. The nature of this relationship, the wage curve, is almost identical across the countries of the world. It is also present, within nations, across different periods of time. Empirical regularity here, as so often, leads to conceptual points. Although it cannot be stressed too forcibly that matters are far from decided, there is some reason to believe that the existence of this curve throws into doubt the relevance of the textbook competitive model of the labor market. By contrast, the curve seems to be like that predicted by a newly emerging kind of macroeconomic theory. Whatever the correct interpretation of the wage curve, it fuses elements from, and seems to have implications for, macroeconomics, labor economics, and regional economics.

International microeconomic data -- comparable random samples from many nations -- offer today's economists something that was out of the reach of Marshall and Hicks. The reason for these surveys' availability is not that the current generation of economists is more intelligent, diligent, or far-sighted than earlier ones, but rather that a new breed of computer designers and government statisticians has changed the raw material with which social scientists can work. Largely unnoticed, a transformation, led by these sorts of innovations, is quietly ticking away inside economics. Research of the kind described in the paper probably only begins to scratch the surface of what might one day be achieved with the test-bed of international survey data. What you see here is perhaps just a taste of the twenty-first century.

Figure 1.1

The wage curve

Table 1 Wage Curves in Twelve Nations

Country	Dependent Variable	Data Set	Coefficient on Log U	t-stat	Fixed effects	Sample size
1. USA	Annual earnings	Current Population Survey 1963-90	-.10	Up to 25	Yes	1730175
2. Britain	Monthly earnings	General Household Survey 1973-90	-.08	6.23	Yes	175500
3. Canada	Gross annual earnings	Survey of Consumer Finances, 1972-87	-.09	6.1	Yes	82739
4. S. Korea	Gross monthly earnings	Occupational Wage Surveys 1971-86	-.04*	25.7	Yes	1359387
5. Austria	Gross monthly earnings	ISSP, 1986&89	-.09	1.59	Some	1587
6. Italy	Gross monthly earnings	ISSP, 1986-9	-.10	0.63	Yes	1041
7. Holland	Net monthly earnings	ISSP, 1988-91	-.17	2.35	Some	1867
8. Switzerland	Net monthly earnings	ISSP, 1987	-.12	3.6	No	645
9. Norway	Gross yearly earnings	ISSP, 1989-91	-.08	2.19	Some	2599
10. S. Ireland	Net monthly earnings	ISSP 1988-91	-.36	1.92	No	1363
11. Australia	Weekly income	IDS 1986	-.19	5.8	No	8429
12. Germany	Gross monthly earnings	ISSP, 1986-91	-.13	1.75	Yes	4629

Note: These are coefficients on local unemployment variables in microeconomic wage equations. The equations are estimated using the numbers of observations stated in the final column.

ISSP is the International Social Survey Programme.

Log U is defined as the logarithm of an area unemployment rate at various levels of disaggregation in different countries. Where indicated by a *, unemployment is measured at the industry level.

The dependent variable, pay, is in natural logarithms. 'Fixed effects' refers to controls for regions or industries. In all equations, personal variables (gender, race, age, schooling, etc.) are included as controls.

Southern Ireland's results are unstable and should be treated with caution.

Where few years of data are available, countries' unemployment coefficients sometimes weaken when full regional dummies are added. This Table is meant only to be representative. Fuller results are in the book The Wage Curve.

References

- Akerlof, George A. and Yellen, Janet L. (1990). 'The Fair Wage-Effort Hypothesis and Unemployment', Quarterly Journal of Economics, 105, 255-284.
- Bhalotra, Sonia (1993). 'Geographical Differences in Unemployment and Wage Rates in India', mimeo, Wolfson College, Oxford University.
- Blackaby, David H. and Manning, D. Neil (1987). 'Regional Earnings Revisited', Manchester School, 55, 158-183.
- Blackaby, David H. and Manning, D. Neil (1990). 'The North-South Divide: Questions of Existence and Stability', Economic Journal, 100, 510-527.
- Blanchflower, David G. and Oswald, Andrew J. (1990). 'The Wage Curve', Scandinavian Journal of Economics, 92, 215-235.
- Blanchflower, David G. and Oswald, Andrew J. (1994). The Wage Curve, MIT Press, forthcoming.
- Card, David (1990). 'Unexpected Inflation, Real Wages and Employment Determination in Union Contracts', American Economic Review, 80, 669-688.
- Christofides, Louis N. and Oswald, Andrew J. (1992). 'Real Wage Determination and Rent-Sharing in Collective Bargaining Agreements', Quarterly Journal of Economics, 65, 985-1002.
- Edin, Per-Anders and Holmlund, Bertil (1989). 'The Unemployment Elasticity of Pay: Evidence from Swedish Micro Data', mimeo, University of Uppsala, Sweden.
- Edin, Per-Anders, Holmlund, Bertil and Ostros, Thomas (1993). 'Wage Behavior and Labor Market Programs in Sweden: Evidence from Micro Data', Working Paper 93-1, University of Uppsala, Sweden, forthcoming in Labour Markets and Economic Performance: Europe, Japan and the US, edited by T. Tachibanaki, Macmillan.
- Freeman, Richard B. (1988). 'Evaluating the European View that the United States has No Unemployment Problem', American Economic Review, Papers and Proceedings, 78, 294-299.
- Freeman, Richard B. (1991). 'Employment and Earnings of Disadvantaged Young Men in a Labor -Shortage Economy', in Pado Schioppa, Fiorella (Ed.) Mismatch and Labor Mobility, Cambridge University Press: Cambridge.
- Greenwald, Bruce and Stiglitz, Joseph E. (1993). 'New and Old Keynesians', Journal of Economic Perspectives, 7, 23-44.
- Groot, Wim, Mekkelholt, E. and Oosterbeek, H. (1992). 'Further Evidence on the Wage Curve', Economics Letters, 38, 355-359.
- Hall, Robert E. (1970). 'Why is the Unemployment Rate So High at Full Employment?', Brookings Papers on Economic Activity, 3, 369-402.
- Hall, Robert E. (1972). 'Turnover in the Labor Force', Brookings Papers on Economic Activity,

3, 709-756.

- Harris, John R. and Todaro, Michael P. (1970). 'Migration, Unemployment and Development: A Two-Sector Analysis', American Economic Review, 60, 126-142.
- Hoddinott, John (1993). 'Wages and Unemployment in Urban Cote d'Ivoire', mimeo, Centre for the Study of African Economies, Oxford University.
- Holmlund, Bertil and Skedinger, Per (1990). 'Wage Bargaining and Wage Drift: Evidence From the Swedish Wood Industry', in Calmfors, L. (ed.) (1990), Wage Formation and Macroeconomic Policy in the Nordic Countries, SNS Forslag, Oxford University Press: Oxford.
- Jowell, Roger, Witherspoon, S. and Brook, L. (1989). (eds.), British Social Attitudes: Special International Report, Gower Press: Aldershot.
- Katz, Lawrence F. and Krueger, Alan B. (1991). 'Changes in the Structure of Wages in the Public and Private Sectors', Research in Labor Economics, 12, 137-172, Jai Press: Greenwich, Conn..
- Katz, Lawrence F. and Krueger, Alan B. (1992). 'The Effect of the Minimum Wage on the Fast Food Industry', Industrial and Labour Relations Review, 46, 6-21.
- Layard, P. Richard G. and Nickell, Stephen J. (1986). 'Unemployment in Britain', Economica, Supplement, 53, 121-170.
- Layard, P. Richard G., Nickell, Stephen J. and Jackman, Richard (1991). Unemployment: Macroeconomic Performance and the Labour Market, Oxford University Press: Oxford.
- Lemieux, Thomas (1994). 'Unions and Wage Inequality in Canada and the United States', in D. Card and R. Freeman (eds.), Small Differences that Matter: Labor Markets and Income Maintenance in Canada and the United States, University of Chicago Press and NBER, forthcoming.
- Lindbeck, Assar (1993). Unemployment and Macroeconomics, MIT Press: Cambridge, Mass.
- Manning, Alan (1993). 'Wage Bargaining and the Phillips Curve: The Identification and Specification of Aggregate Wage Equations', Economic Journal, 103, 98-118.
- Montgomery, Edward B. (1993). 'Patterns in Regional Labor Market Adjustment: The US Versus Japan', NBER Paper number 4414.
- Moulton, Brent R. (1986). 'Random Group Effects and the Precision of Regression Estimates', Journal of Econometrics, 32, 385-397.
- Oswald, Andrew J. (1986). 'Wage Determination and Recession: A Report on Recent Work', British Journal of Industrial Relations, 24, 181-194.
- Phelps, Edmund S. (1990). 'Effects of Productivity, Total Domestic Product Demand and Incentive Wages on Unemployment in a Non-Monetary Customer-Market Model of the Small Open Economy', Scandinavian Journal of Economics, 92, 353-368.
- Phelps, Edmund S. (1992). 'Consumer Demand and Equilibrium Unemployment in a Working Model of the Customer-Market Incentive-Wage Economy', Quarterly Journal of

Economics, 107, 1003-1032.

Phillips, A. William (1958). 'The Relation Between Unemployment and the Rate of Change of Money Wage Rates in The United Kingdom, 1861-1957', Economica, 25, 283-299.

Rebick, Mark (1994). 'Widening Firm Size Differentials in Japan', Journal of the Japanese and International Economies, forthcoming.

Rowthorn, Robert E. (1977). 'Conflict, Inflation and Money', Cambridge Journal of Economics, 1, 215-239.

Sargan, J. Dennis (1964). 'Wages and Prices in the United Kingdom: A Study in Econometric Methodology', reprinted in Hendry, David F. and Wallis, Kenneth F. (eds.) (1984). Econometrics and Quantitative Economics, Basil Blackwell: Oxford.

Shapiro, Carl and Stiglitz, Joseph E. (1984). 'Equilibrium Unemployment as a Worker Discipline Device', American Economic Review, 74, 433-444.

Solow, Robert M. (1990). The Labor Market as a Social Institution, Basil Blackwell: Oxford.

Wagner, Joachim (1994). 'German Wage Curves: 1979-1990', Economics Letters, forthcoming.

Winter-Ebmer, Rudolph (1994). 'Unemployment and Individual Pay: Wage Curve or Compensating Differentials?', Labour Economics, forthcoming.

Endnotes

¹ This includes Bhalotra (1993), Blackaby and Manning (1987, 1990), Card (1990), Christofides and Oswald (1992), Edin and Holmlund (1989), Edin et al (1993), Freeman (1988, 1991), Groot et al (1992), Hoddinott (1993), Holmlund and Skedinger (1990), Katz and Krueger (1991, 1992), Lemieux (1994), Montgomery (1993), Rebick (1994), Wagner (1994), and Winter-Ebmer (1994). Some preliminary findings for Britain and the U.S. are reported in Blanchflower and Oswald (1990).

² However, there is evidence that the unemployment elasticity of pay differs for different sub-groups. It is higher for the young, the unskilled, and those outside unions.

³ An interesting modern defence of Phillips Curve methods can be found in Manning (1993).