

A review of “Unemployment. Macroeconomic Performance and the Labour Market”, by Richard Layard, Stephen Nickell, and Richard Jackman

Olivier Blanchard*

July 27, 2006

“Unemployment. Macroeconomic Performance, and the Labour Market”, by Richard Layard, Stephen Nickell, and Richard Jackman, (LNJ in what follows) is the second edition of a book first published in 1991. The second edition is identical to the first except for a long introduction, which reviews how the conclusions reached in the first edition have withstood the test of time.

1. The book

The book was and remains an impressive achievement. The way to read it however is not so much as a treatise than as a manual of battlefield surgery. Its purpose is clear: How to understand, and then reduce unemployment in Europe, by taking inventory of the knowledge at hand.

To understand the authors’ state of mind, it is important to put the book in context. Starting in the late 1970s, and continuing throughout the 1980s, unemployment steadily increased in Europe. Just as economists had converged on the idea of a natural unemployment rate, this natural rate appeared to move over

* MIT and NBER.

time, in dramatic and largely mysterious ways. Unemployment became the main economic issue facing Europe and dominating economic policy discussions.

This led the LNJ team to jump into action, and for most of the 1980s, they were at the forefront of a large research effort on unemployment. Together with others, they developed a general but loose framework to organize thoughts on the issue. The framework went something like this:

Forget the conventional analysis of the labor market in terms of neo-classical labor demand and supply: Unemployment does not make sense in that framework, and imperfections must be at the center of any theory of unemployment. Think instead in terms of a price-setting and of a wage-setting relation.

Think of the price-setting relation as giving the price chosen by imperfectly competitive firms. As the price-setting relation gives the price given the nominal wage (among other factors), it implicitly gives us a first relation between the real wage and employment.

Now think of the wage-setting relation as giving the nominal wage set in the labor market. As the wage-setting relation gives the nominal wage given the price level (among other factors), it gives us a second relation between the real wage and employment (or unemployment).

The natural rate of unemployment is then the unemployment rate which, if we leave aside nominal rigidities, reconciles the real wage implied by price-setting and the real wage implied by wage-setting. In the presence of nominal rigidities, the actual unemployment rate may differ from the natural unemployment rate. Deviations of the actual from the natural rate will then lead to movements in price and wage inflation.

In the end, the price-setting and wage-setting relations resemble the old labor demand and labor supply: In the presence of decreasing returns to labor, price setting implies an increase in the price given the wage as employment increases. Put another way, it implies a negative relation between employment and the real wage—just like labor demand. Given the price level, wage setting implies higher nominal wages if unemployment is low—equivalently if employment is high. Put another way, it implies a positive relation between the real wage and

employment—just like labor supply. Still, a lot is gained from the conceptual switch: Thinking in terms of wage and price setting suggests a large set of potential determinants of the natural rate of unemployment, from labor market institutions such as unemployment insurance or employment protection, to product market characteristics such as the degree and scope of regulation. It also implies that the effects of changes in the economic environment—from changes in the user cost of capital to changes in the rate of productivity growth—on unemployment may now depend on the labor and product market institutions themselves. In short, it gives a rich organizing framework for thinking about and looking at the evidence on the relation of unemployment to institutions and shocks.

This framework is presented in Chapter 1 of LNJ (the first chapter of the first edition, now preceded by the new introduction), an amazing tour de force which shows, in 82 pages, how the framework can be developed, estimated, and used to answer all the questions you may have had about unemployment. For those readers who do not intend to do research on unemployment, reading Chapter 1 may be all they need to do to get what they need from the book.

Still, without a lot more detail in what lies behind price and wage-setting, the framework remains a shell. So, the next and major part of the book, namely Chapters 2 to 7, looks more closely at the various theories of price and wage setting.

In the case of wage setting, we (the profession) suffer from an overabundance of theories: Some emphasize the role of collective bargaining, with a focus on the objective function of unions and on the level of bargaining. Some, generically known as “efficiency wage theories” emphasize the complexity of labor relations within the firm, the need for firms to maintain morale, to retain workers. Yet others, called “search models” or “matching/bargaining models”, focus on the bargaining power that emerges from frictions in the labor market. Unfortunately—for those of us who would like the world to be simple—all three strands are relevant. In Chapters 2 to 6, LNJ provide a detailed presentation of each of the strands, with no attempt however at integration. (Given what was known then, and what is known now, I find the presentation of search models to be weaker

than the others. More on this below). In each case, the presentation of theory and facts is exhaustive, indeed sometimes exhausting. No theory gets away without a thorough empirical confrontation, and a few do not survive the confrontation.

In the case of price-setting, the profession's thinking is much more primitive: Clearly, many firms have enough market power that it makes sense to think of them as choosing prices. How the markup they choose is determined and how it varies as a function of the economic environment remains largely a mystery. In Chapter 7, LNJ take stock of what we know and we do not know.

Having surveyed potential micro foundations, LNJ then return in Chapters 8 and 9 to the macro implications and the macro evidence. They first take the price-setting and wage-setting model through its paces, in both a closed and an open economy. They then take the empirical plunge, looking at the relation between unemployment, shocks, and measures of institutions, across time, and across countries.

They estimate price-setting and wage-setting equations for each OECD country, and then regress the estimated coefficients on measures of labor market institutions. The various theories they have reviewed are used loosely, mostly as a guide to which variables should be put on the right-hand side of each equation. Separate identification of the price and wage equations is a difficult task however, and LNJ's approach to identification is not particularly convincing. Probably for that reason, they also explore a reduced form approach, in the form of a cross-country regression of unemployment directly on the measures of institutions—an approach which is valid under the maintained assumption that the different countries have experienced the same underlying shocks.

The book culminates and ends with Chapter 10, a chapter on policy. The recommendations are unambiguous: Yes to generous unemployment benefits, so long as they are of limited duration. Yes to carrots but also to sticks to get the unemployed into jobs if such jobs are available. Yes to either centralized or very decentralized bargaining, but not to bargaining structures in between. Yes, in some cases, to tax based incomes policies. No to public employment. No to work sharing. No to early retirement. By the end of the chapter, the manual is com-

plete, the recommendations sharp and clear; now it is up to policy makers to make things happen.

As my description must have made clear, this is not a standard academic tome. A sense of urgency permeates the whole book. It starts with the epigraph, “To the millions who suffer through want of work”. It continues with a continual back and forth between theories and data. The treatment eschews elegance and simplicity in favor of eclecticism and, sometimes, messiness. LNJ clearly do not believe that one theory will fit all, and are unwilling to leave aside any insight or policy that may help improve things. Yet, eclecticism does not lead them to equivocate. They conclude that enough is known to give advice. The last chapter reads more like the platform of a (remarkably intelligent) political party than a guarded academic assessment of the state of knowledge on unemployment.

In that sense, the book is very different from the other two important books on unemployment published in the early 1990s:

The first was written by Edmund Phelps in 1994. Many of the themes are similar to those of LNJ. Like LNJ, Phelps focuses on the role of imperfections in labor and product markets. But, unlike LNJ, he chooses to explore the implications of specific imperfections, namely a particular form of efficiency wages for the labor market, and firms investing in a customer base in the goods market. The result is a much tighter theoretical framework, but a more constrained look at the evidence. (Phelps was actually the reviewer of the first edition of LNJ in this journal, in 1992. The review, which is as much about Phelps’ own work than about LNJ, shows the sharp differences in research strategy between the two books. Phelps rightly points to the holes in LNJ’s broad theory, for example in the way productivity growth affects unemployment—a theme I shall return to below. But the tight focus of Phelps’ book has its downside: The particular choice of imperfections is not universally shared, and Phelps’ book has had much less impact than LNJ’s.)

The second was written by Pissarides in 1990 (with a second edition in 2001). While LNJ were working in the 1980s on their synthetic approach, Pissarides was exploring how the earlier work by Diamond and by Mortensen on search

markets could shed light on unemployment. His 1990 book was a summary of his research to date, and a systematic exploration of the implications of the matching/bargaining model for unemployment. The theory is elegant, and, as I shall discuss below, has been enormously influential. But again, the book serves a different purpose from LNJ, namely the systematic exploration of a particular line of theory, rather than the writing of a manual for applied economists and policy makers.

2. Facts and research since 1991

A lot has happened since 1991, both with respect to unemployment, and with respect to research on unemployment (For a more chronological description of events and ideas since the early 1970s, as well as a detailed bibliography, see Blanchard [2006]):

Unemployment has remained very high in continental Europe, although with increasing differences across countries. The large countries, Germany, France, Italy, and Spain still have unemployment rates close to or above 10%. Many small countries, among them the Netherlands and Denmark, have low unemployment.

On the theory front, work has continued unabated, although progress has been uneven. Little progress has been made on wage setting under collective bargaining, and on nailing down the relevant objective function of unions. Little progress has also been made on wage setting with efficiency wages, and on identifying the respective roles of morale, retention, and other motives that firms may have in setting wages. Much progress has taken place however in the exploration of matching/bargaining models. A benchmark model has emerged (often referred to as the Diamond-Mortensen-Pissarides (DMP) model, and presented in the second edition of the book by Pissarides), and has been used as a point of departure for many extensions. (A useful review of some of these extensions was given recently in this Journal by Rogerson et al [2006]).

It is easy to see why this line of modeling has become so popular. While abstract, the model appears to fit many of the aspects of actual labor markets: Workers become unemployed either by choice, but more often by accident. The unemployed

search for jobs. They choose how hard to search, what reservation wage to hold out for. Workers have some bargaining power, coming from the cost to firms of firing them and finding adequate replacements. As a result of this easy mapping between theory and facts, components of the model, such as for example the search behavior of the unemployed, can easily be confronted to micro data. And, in turn, findings from micro studies of the effects of policy changes can be taken back to the model, to allow for a better sense of the relevant parameters.

At the same time, it is also clear that this line of modeling will not do the job by itself, and that the two other lines taken up by LNJ, namely collective bargaining and efficiency wages, are very relevant. First, collective bargaining remains, at least in continental Europe, essential to wage determination, and, I believe, essential to an understanding of differences in unemployment evolutions across European countries. Second, the formalization of bargaining in the standard DMP model through Nash bargaining, while elegant and powerful, is a very poor description of reality. As a number of researchers have pointed out, real wages appear to move much less than is implied by Nash bargaining. Constraints coming from intra-firm efficiency-wage considerations look like plausible candidates to explain these wage rigidities.

On the empirical front, the increasing availability of micro data on individuals and firms and the development of employer-employee matched data sets have led to an explosion of careful micro-empirical work.

At a descriptive level, household surveys, developed along the lines of the US Current Population Surveys, have given us a characterization of the flows of workers between employment, unemployment, and inactivity, across types as well as over time. Firm surveys have given us a characterization of job flows, again across firm types and over time (one must single out here the work of Davis and Haltiwanger [1996], which has been highly influential).

At the same time, the many attempts by European governments to reduce unemployment have given us a large number of quasi-natural experiments, and a better sense of the effects of labor market institutions on the labor market. What we have learned the most about is unemployment insurance. The evidence is that

limiting the duration of benefits, as well as making them more contingent on job search and job acceptance, leads to more active search, a lower reservation wage, and lower duration of unemployment (see for example Frederiksson and Holmlund [2003] for a survey). We have learned less about the effects of employment protection. We know that, across countries, higher employment protection is associated with longer unemployment duration, lower participation, and lower flows in and out of unemployment. Convincingly establishing causality has been more difficult however: Clean experiments, and appropriate data, are harder to get.

In parallel, research has continued to explore, at the macroeconomic level, the relation of unemployment to shocks and to institutions. In effect, it has extended the initial cross-country reduced form approach of LNJ to panel data, using better measures of institutions and shocks. Having participated in the effort (Blanchard and Wolfers [2000]), I feel freer to criticize it. I see these panel data regressions as useful data descriptions. As such, they suggest a relatively robust set of correlations between unemployment, some adverse shocks, and some labor market institutions—in particular unemployment benefit duration, and the degree of coordination in bargaining (how much institutions affect unemployment directly and how much they affect unemployment by mediating the effects of shocks has proven harder to establish.) However, asking these panel data regressions to tell us conclusively about causal effects of institutions, shocks, and interactions of shocks and institutions, on unemployment is beyond what they can deliver. Causality is next to impossible to establish, as many institutional changes are triggered by labor market developments. And despite a remarkable effort, in particular by the OECD, to construct quantitative measures of institutions across countries over time, it is not obvious that one can really construct appropriate indexes for such multi-dimensional institutions as collective bargaining or employment protection (I return to this issue below).

3. How much do we know?

Let me take up two questions here. First, has the book withstood the test of time (or at least the test of 15 more years)? Second, do we (the profession) have a

good understanding of what determines unemployment, and a good explanation of the variations in the unemployment rate across European countries, and over time?

As to the answer to the first question, the authors give themselves a satisfecit. In their new introduction, they state: “We conclude that our previous analysis performs well, which is why we have responded positively to requests to republish the book”. I basically agree. The general framework has proven very useful. Indeed, it has made its way into textbooks, as it allows us to link the derivation of the natural rate of unemployment to the “aggregate supply relation” between activity and prices. While theory has advanced since 1991, in particular with respect to search models, the book remains a very useful introduction to the different models of wage and price determination. And the largely consensual list of recommendations that economists would make to policy makers today would be close to the list given in Chapter 10: Generous unemployment benefits, but more explicitly contingent on job search and job acceptance. Some employment protection, perhaps in the form of experience rating, but with limited judicial intervention. A limit to payroll taxes (and by implication to the benefits that these taxes finance). And no to Malthusian solutions, from mandatory work sharing to early retirement. And indeed, while progress has been slow, most European governments are moving in that direction. A rough but intriguing exercise reported in the new introduction is a cross-country regression of the change in the unemployment rate since 1980 on the number of employment-friendly and the number of employment-unfriendly institutional changes over the same period. The signs on both right-hand-side variables are correct, and the fit is quite good. Maybe LNJ and the rest of us have been giving the right advice.

As to the answer to the second question, I am less sanguine. To state an old adage, we know a lot, but there is a lot we do not know. Here is a list of four dimensions in which I wish I knew more.

How productivity gets into wages

Increases in productivity growth appear to be reflected over time in equal in-

creases in wages. But what is the mechanism through which productivity gets into wages? There are at least three different ways of thinking about it.

The first, which interprets all non-work as leisure, relies on the income effect: Productivity increases consumption, which decreases the marginal utility of consumption and increases the reservation wage. Indeed, if preferences are such that productivity growth is associated with a balanced growth path and a constant employment rate along that path, and if consumption increases in proportion to productivity, the reservation wage increases instantaneously and in proportion to productivity, without any change in employment. (If consumption increases in the short run less than one-for-one with productivity, then an increase in productivity will have an effect on employment for some time. This is the mechanism central to real business cycle models.)

The second, which distinguishes unemployment from leisure, assumes that unemployment provides no direct utility, and thus the reservation wage depends only on the level of income if unemployed. Then the adjustment of the reservation wage to productivity depends on the adjustment of unemployment benefits and of non-labor income to productivity (the second channel is the channel emphasized by Phelps in his book). If unemployment benefits and non-labor income eventually adjust to productivity, so will the reservation wage.

The third, which is more relevant when collective bargaining is prevalent, thinks of wage agreements as setting wage growth based on “normal” productivity growth—presumably some distributed lag of productivity growth. To the extent that normal productivity growth eventually adjusts to actual growth, so will the growth of reservation wages.

All three approaches deliver, to a first order, long run neutrality of unemployment to productivity growth. But they have very different short and medium run implications. And, given the large low-frequency movements in productivity growth—down in the 1970s and 1980s, up in the United States since the mid 1990s—these differences matter very much. The issue arises in other contexts as well. When firms respond to high labor costs by increasing capital intensity and thus by increasing productivity growth above the underlying rate of technological

progress, do workers take into account the difference between productivity growth and underlying progress, or do they respond to higher productivity growth with yet higher wage demands? The truth is that we do not know. And this is a serious hole in our knowledge.

Employment protection, flows, and turbulence

Thanks to the availability of panel data on firms, we have a much better understanding of the relation between unemployment, reallocation, and productivity growth. But from these data have also emerged two puzzles:

First, given the widely held perceptions of a sclerotic Europe and a dynamic United States, one would have expected lower job flows and reallocation in Europe than in the United States. This does not appear to be the case: The level of job flows appears rather similar on both sides of the Atlantic. This fact may hide a different composition of separations, with a lower proportion of efficient separations and a higher proportion of inefficient separations in Europe. But, as it stands, it forces us to be more careful about the role of labor market institutions, in particular employment protection.

Second, given the widely held perception of increased reallocation from both goods market deregulation and from globalization, one would have expected reallocation to increase over time. The issue was taken up in LNJ who showed that, surprisingly, rough measures of reallocation intensity, such as the standard deviation of employment growth across sectors, did not appear to have increased over time. We now have much better measures of reallocation, namely job flows, as well as longer time series. And, surprisingly, these measures give the same signal as the earlier and cruder measures: Job flows do not appear to have significantly increased. One may conclude that increased turbulence has taken other forms, but, as it stands, I also see this as a serious puzzle, one we need to look at more closely.

Collective bargaining and trust

Today, there is fairly wide agreement about the optimal design of such labor market institutions as unemployment insurance, employment protection, or the use of a negative income tax or the minimum wage. Some countries appear to have gone further than others in the right direction, and, in the process, have achieved low unemployment. Until recently, the “Dutch model,” was presented as the model to follow. The love affair has now turned to Denmark, and its system of “flexsecurity”. While other continental European countries have indeed something to learn from these countries, the question, simply put, is the following: If France or Italy were to import the Dutch or Danish systems of unemployment insurance and employment protection, would they converge to Dutch or Danish natural unemployment rates?

I am skeptical that this would be the case. The reason is that the functioning of European labor markets appears to depend on more factors than just these formal institutions. It appears to depend on the nature of collective bargaining, and, more generally and more fuzzily, on the degree of trust between social partners—labor, firms, and the state. Countries with high reported trust between labor and firms have had a lower increase in unemployment than others, and the correlation appears to reflect in part causality. The dramatic turnaround of the United Kingdom, a country with notoriously low trust between labor and firms until recently, suggests that trust should not necessarily be taken as a country fixed effect. Thus, understanding both where trust comes from and how it affects outcomes, should be high on our agenda.

Back to European unemployment

A major theme of LNJ is that the evolution of the unemployment rate is shaped by both shocks and institutions. The broad historical description they gave in 1991 blamed adverse shocks—in particular the two increases in the price of oil in the 1970s—for the initial increase in the natural rate of unemployment, and labor market institutions for the persistence of this high natural rate. These institutions, they argued, had led to a large increase in the duration of unemployment in response to the adverse shocks; the near irrelevance of the long-term unemployed in wage setting implied little downward pressure on wages and on unemployment.

This interpretation was challenged by Phelps in his JEL review: It seemed improbable that these initial shocks, which indeed had been largely reversed later, could still be responsible for high unemployment more than ten years later. And, indeed, another fifteen years later, the notion that these initial shocks would still have much relevance in explaining unemployment today, seems even less plausible. This has led to a two-fold search:

First, a search for other adverse shocks. Phelps himself emphasized factors that increased the real rate of interest, and through the interest rate, increased the natural rate of unemployment. Others have suggested shifts in the relative demand for labor together with downward constraints on unskilled wages. Yet others have suggested the importance of an increase in turbulence—as we have seen, so far with mixed success.

Second, a search for changes in institutions. A plausible story is that, in response to the initial increase in unemployment, governments reacted by taking the wrong measures. To alleviate the pain of unemployment, they increased the generosity and duration of benefits. To limit the increase in unemployment, they tried to prevent firms from laying off workers, through tougher employment protection. To better share the burden of low employment, they used early retirements, and work sharing. All these measures then in turn increased the natural rate, even as the initial shocks disappeared. The story is plausible, and I believe, largely right (and it is the story that Stephen Nickell would tell today. See Nickell et al [2005]). The problem with it is that available measures of labor market institutions do not show the sharp deterioration of labor market institutions implied by the story. Based on my study of the history of French institutions, I believe that the problem is with the crude measures of institutions we have developed, not with the story. But this, and by implication the effect of labor market institutions on unemployment, is yet another issue we need to resolve before we can declare intellectual victory.

References

- [1] Olivier Blanchard. European unemployment: The evolution of facts and ideas. *Economic Policy*, 1:1–54, January 2006.
- [2] Olivier Blanchard and Justin Wolfers. Shocks and institutions and the rise of European unemployment. The aggregate evidence. *Economic Journal*, 110(1):1–33, March 2000.
- [3] Steven Davis, John Haltiwanger, and Scott Schuh. *Job Creation and Job Destruction*. MIT Press; Cambridge, Mass, 1996.
- [4] Peter Fredriksson and Bertil Holmlund. Improving incentives in unemployment insurance; A review of recent research. *CESifo working paper 922*, April 2003.
- [5] Stephen Nickell, Luca Nunziata, and Wolfgang Ochel. Unemployment in the OECD since the 1960s: What do we know? *Economic Journal*, 115:1–27, January 2005.
- [6] Edmund Phelps. Customer demand and equilibrium unemployment in a working model of the incentive-wage customer-market economy. *Quarterly Journal of Economics*, pages 1003–1033, August 1992.
- [7] Edmund Phelps. *Structural Slumps. The modern equilibrium theory of unemployment, interest, and assets*. Harvard University Press; Cambridge MA, 1994.
- [8] Christopher Pissarides. *Equilibrium Unemployment Theory*. Basil Blackwell; Oxford, 1990.
- [9] Christopher Pissarides. *Equilibrium Unemployment Theory*. Second edition, MIT Press, 2000.
- [10] Richard Rogerson, Robert Shimer, and Randall Wright. Search-theoretic models of the labor market: A survey. *Journal of Economic Literature*, 43(4):959–988, December 2005.