Lecture 1. Shocks, factor prices, and unemployment.

Olivier Blanchard*

October 2000

A simple–minded approach to unemployment is to start with a labor demand curve and a wage setting curve (starting with a standard labor supply would be *too* simple–minded, at least if the goal is to explain the evolution of European unemployment), look at the equilibrium, and trace the effects of various shifts on equilibrium wages and unemployment.

This is simple–minded in at least three ways. First, it ignores the role of nominal rigidities and fluctuations in aggregate demand which are clearly central to short–run fluctuations in output and unemployment. Second, the wage setting curve depends very much on the nature of imperfections and institutions in the labor market. If one wants to understand how it looks and why it shifts, one must look more deeply into its origins. And, third, many of these imperfections and institutions imply that firms may not be operating on their labor demand curve. In many of the models of bargaining for example, the wage is typically not equal to the marginal product of labor.

Yet, and this is the theme of this lecture, this simple–minded approach is very useful. It makes straightforward and strong predictions: In the medium run, if factor prices are too high—be it the price of labor or the

*Lecture 1, Lionel Robbins Lectures, LSE, to be given in October 2000. Preliminary. Comments welcome.

price of capital—unemployment will increase. And, if unemployment is too high, the only way to decrease it, again in the medium run, is to decrease one or both factor prices. These predictions are likely to remain true in more complex models (I shall return to this issue in the second lecture). And they provide a very useful guide to the evolution of European unemployment, both on the way up, and more recently, on the way down.

I start the lecture by sketching the model. I take it through its paces, by looking first at the effect of a decrease in the rate of technological progress, then at the effect of a sharp but temporary rise in interest rates as might arise from a monetary contraction. I then turn to the data, returning to the role of the productivity slowdown of the 1970s, to the role of monetary policy in shaping equilibrium unemployment in the 1970s and 1980s, and finally to the role and the effects of wage moderation in the decrease in unemployment in the Netherlands and in Ireland since the early 1980s.

1 Laying out the model.

Think of an economy which is growing along its balanced growth path. Firms use capital and labor under constant returns to scale. To allow for balanced growth, technological progress is assumed to be labor augmenting (equivalently: Harrod neutral), so the production function is given by:

$$y = F(an, k)$$

where y denotes aggregate output, n and k denote aggregate employment and capital, and a denotes the technological level. Put another way, output is produced using two inputs, "labor in efficiency units" an, and capital k. Assume that a grows at the constant rate g_a. I shall refer to g_a as the rate of labor–augmenting technological progress, or the rate of technological progress for short.

Labor demand

Assume that, in the short run, capital is fixed. The competitive, profit–maximizing, demand for labor given capital (short–run labor demand for short) can then be written as:

$$\frac{an}{k} = f(\frac{w}{a}) \qquad f'(.) < 0 \tag{1.1}$$

The ratio of labor in efficiency units to capital is a decreasing function of the wage in efficiency units. For our purposes, it will be more convenient to rewrite the relation as:

$$n = \frac{k}{a} f(\frac{w}{a}) \tag{1.2}$$

This relation between employment, n, and the wage in efficiency units, w/a, is drawn for a given value of k/a as the downward sloping curve DD in Figure 1. Its slope depends on the short–run elasticity of substitution between capital and labor: the lower the elasticity the steeper the short–run demand curve.

Turn to the long run. Assume the user cost of capital is given, and equal to c (think either of a small open economy which takes the interest rate, and thus the user cost of capital as given; or just think of what we do as a partial equilibrium analysis, taking the interest rate as given). In the long run, the wage in efficiency units must be such that the associated profit rate is equal to the user cost. Thus, long run labor demand is given by:

$$c = \pi = g(\frac{w}{a}) \qquad g'(.) < 0 \tag{1.3}$$

where π is the profit rate, and $g(.)$ is the factor price frontier relation implied by the production function $F(.,.)$. In the long run, the wage in efficiency units must be such as to generate a profit equal to the user cost—equivalently, such as to generate zero net profit ($\pi - c = 0$). This relation is

drawn as the horizontal line LL in Figure 1.

Wage setting

Turn next to the wage setting relation. For notational simplicity, it is convenient to normalize the labor force to 1 so u denotes both the unemployment level and the unemployment rate. Given that, along the balanced growth path, both the wage in efficiency units and the unemployment rate are constant, it is natural to write the wage setting relation as:

$$\frac{w}{a} = z\, h(u) \qquad\qquad h'(.) \leq 0 \tag{1.4}$$

This states that the wage (in efficiency units) is a decreasing function of unemployment. The parameter z is a black box parameter, which captures all the factors which may affect the wage given the unemployment rate, from unemployment benefits, to the structure of bargaining, and so on. In this lecture, I shall not open the black box and look further into the determinants of z or the form of the function $h(.)$. But we shall spend much of the next two lectures looking inside the black box.

Assume that, to start with, the economy is growing along its balanced path. The equilibrium is then given by point A in Figure 1. The short–run labor demand, long–run labor demand, and wage setting curves all go through point A. Equilibrium unemployment is equal to u (unemployment is measured from right to left, starting at 1, the labor force.) At the equilibrium unemployment rate u, the wage in efficiency units implied by wage setting is consistent with a profit rate equal to the user cost. Output, capital, and employment in efficiency units all grow at rate g_a.

[Figure 1. The balanced path.]

To see what this model implies, the next two sections take it through its paces, looking first at the effects of a decrease in the rate of technological

Figure 1. The balanced path

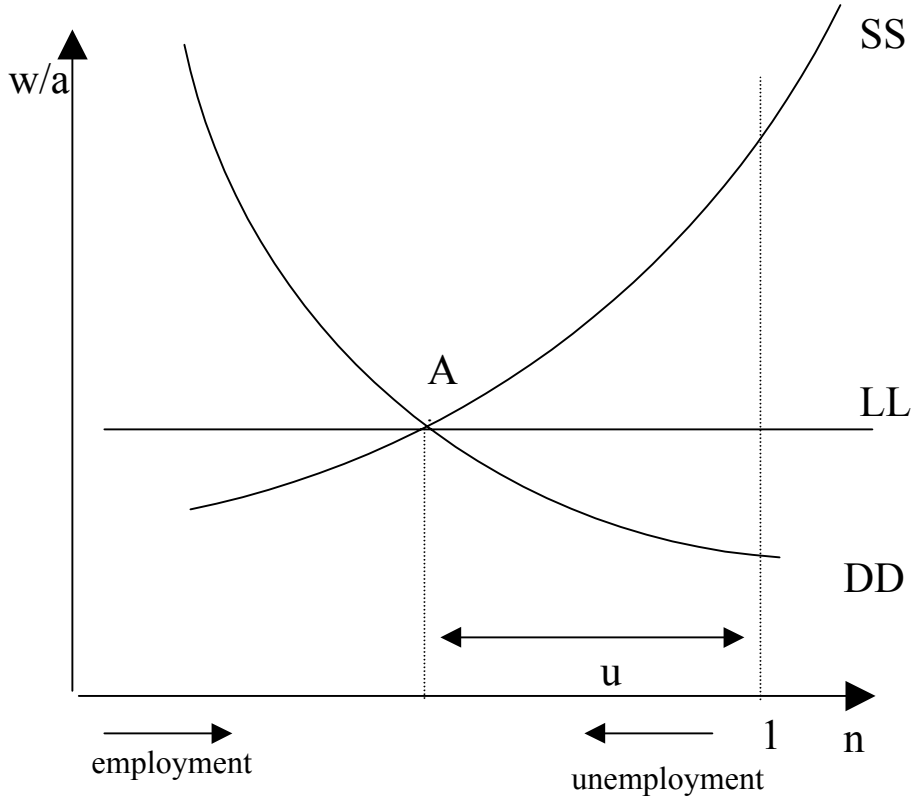

progress, and then at the effects of an increase in real interest rates.

2 The effects of a decrease in the rate of technological progress.

Looking at the five major European countries (Germany, France, the UK, Italy, and Spain) the rate of technological progress, which had been close to 5% in the 1950s and 1960s, decreased to 3% in the first half of the 1970s, and to 2% in the second half of the 1970s.[1] It has remained around 2% since.

It is natural to ask how much of the increase in unemployment might have come from such a decline and how long the effects may have lasted.[2]

The first answer given by the model above is that, if this decrease had been instantaneously perceived and understood by all, it would have had no impact on equilibrium unemployment. Output, capital, employment in efficiency units, and wages, would have grown at a lower pace, the unemployment rate would have remained the same.

But such an exceptional decrease is likely to take some time before it is fully understood by the economic players.[3] Measures of productivity growth move a lot from year to year, and most of the movements are transitory.

[1] These numbers refer to technological progress in the business sector. The rate of labor–augmenting technological progress is obtained by first constructing the Solow residual for each year, and dividing it by the share of labor in the business sector for that year. Put another way, the rate of labor augmenting technological progress is equal to the rate of total factor productivity growth divided by the labor share.

[2] I am obviously not the first one to ask... One of the first systematic attempts was made by Bruno and Sachs [1985], and there have been many since.

[3] Bob Solow, in commenting on this lecture, used the term "comprehension lag", which sounds just right.

Workers, as well as others, are likely to take some time to realize that the change is actually permanent. Meanwhile, they are likely to extrapolate the old trend and thus overestimate the value of a.

A simple way to capture this idea is to extend the wage setting curve to read:

$$\frac{w}{a^*} = zh(u) \tag{2.1}$$

where a^* is the perceived rather than the actual level of technology.[4] Rewriting gives:

$$\frac{w}{a} = z'\, h(u) \qquad \text{where } z' \equiv z\, \frac{a^*}{a}$$

So, after a permanent decrease in g_a, a^* is likely to increase faster than a for some time, leading to an increase in z', and a shift of the wage setting to the left over time, from SS to, say, SS' in Figure 2. As reality sets in, z' starts decreasing, and the wage setting curve starts shifting back to the right. As a^* eventually returns to a, z' goes back to z, and the wage setting relation returns to SS.

What happens to unemployment along the way? As SS initially shifts to the left, w/a increases. And this higher wage has two effects on firms. First, given their capital stock, they reduce their demand for labor. The economy moves along DD. But, also, the higher wage implies a profit rate below the user cost, and thus a decrease in capital accumulation relative to the balanced growth path; over time DD shifts to the left. Employment declines for two reasons: a lower employment given capital, as well as a lower capital

[4] Another interpretation which is often given, and which I find attractive but difficult to make operational, is to think of the underlying rate of technological progress as shaping "wage aspirations." Workers get used to and expect a given rate of wage increase. In the face of a decrease in the underlying feasible rate, it then takes a long time for them to adjust their aspirations to the new reality.

accumulation. The specific dynamics depend on the details of the model, but the general adjustment path can be represented by the counterclockwise loop in Figure 2: a period of higher wages, and higher unemployment until aspirations adjust, wages decline, and the economy recovers.

[Figure 2. A decrease in the rate of technological progress]

That the decline in the rate of technological progress had something to do with the increase in unemployment in the 1970s will not be seen as very controversial.[5] The more interesting but more difficult question is how much and for how long.

Because I could not find an answer in the literature, I decided to do a rough calibration. What follows is very much a back–of–the–envelope computation. The exercise requires two components: First, a way of thinking about the effects of the decline in g_a and the evolution of a^* relative to a. Second, the use of a quantitative model along the lines above, but with more explicit dynamics, to trace the effects of the evolution of (a^*/a) on unemployment.

To think about the first component, assume that g_a is subject to both transitory and (infrequent) permanent movements. Suppose that workers suspect that, from some time, say $t = 0$, there may have been a permanent decrease in g_a. Then, from $t = 0$ on, they will adjust g_a^* according to:

$$g_{at}^* = \lambda g_{at-1}^* + (1 - \lambda) g_{at} \tag{2.2}$$

where λ will depend on the variances of the permanent and transitory

[5] We may be seeing the beginning of the reverse experiment in the United States today. The unexpected and persistent increase in the rate of total factor productivity growth since the mid to late 1990s may well be one of the main factors behind the decrease in the equilibrium rate of unemployment of the last few years.

Figure 2. A decrease in the rate of technological progress

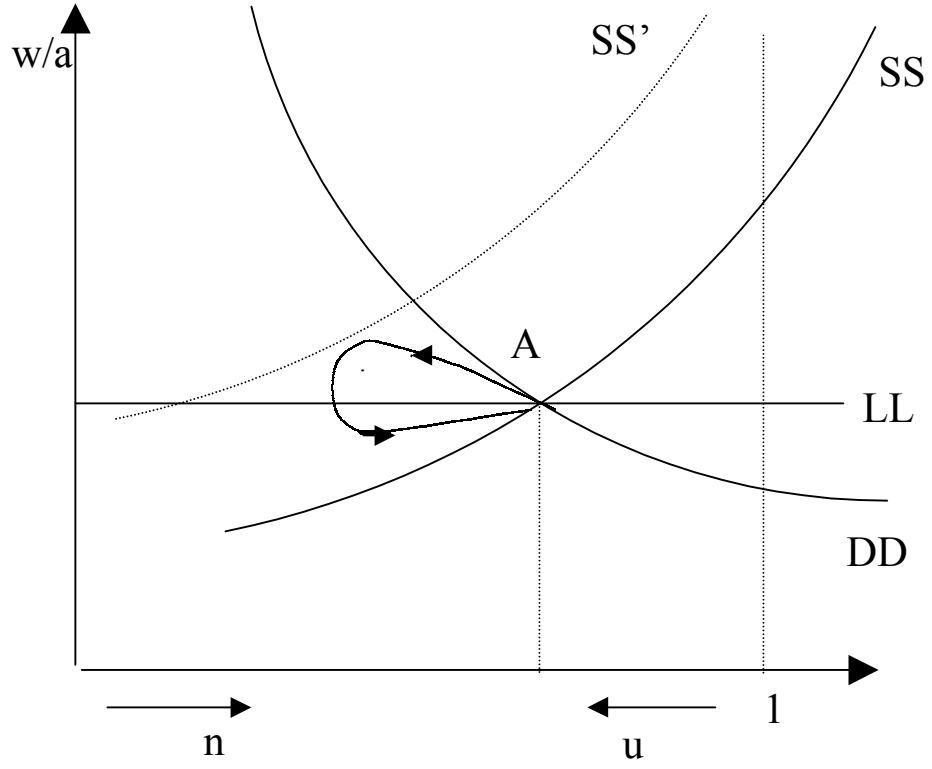

shocks to g_a. And, given their perceived growth rate, they will then compute their perceived level of a according to:

$$\log a_t^* = \log a_0 + t g_a^*$$

As workers revise down their estimated underlying rate of technological progress, g_a^* will steadily converge to g_a. But as workers use too high a growth rate to compute a^*, a^* will diverge from a for a while, until eventually returning to a over time.

To get a sense of magnitudes, suppose that at time $t = 0$, g_a decreases from 5 to 2%. Figure 3 then shows the evolution of $\log a$ and $\log a^*$ over time if $\lambda = 0.9$. In this case, $\log a^* - \log a$ reaches a maximum of 10% roughly 10 years after the start of the slowdown, and then returns to zero over time. The point of this exercise is simple: The effect of the decrease in the rate of technological progress can be fairly large (equivalent to a 10% "wage push," an increase of z of 10%), and the maximum effect may happen after some time, here after 10 years.

[Figure 3. Actual and perceived tfp levels.]

The second component we need is a quantitative model along the lines of the model sketched above. I developed such a model earlier (Blanchard [1997], and Blanchard [1998]), and I shall rely on it here. (The appendix gives the specific equations).

The model formalizes firms as competitive profit maximizers, with a CES production function, and labor–augmenting technological progress. The simulations below assume an elasticity of substitution of 1 (some simulations in Lecture 2 will relax this assumption.)

There are two sources of dynamics, coming from costs of adjustment for both capital and for factor proportions.

- The lower the costs of adjustment to capital, the faster the difference

Figure 3. Actual and perceived technology levels. 1960= 1.0

Effects of a decrease in g_a from 5% to 2% in 1970

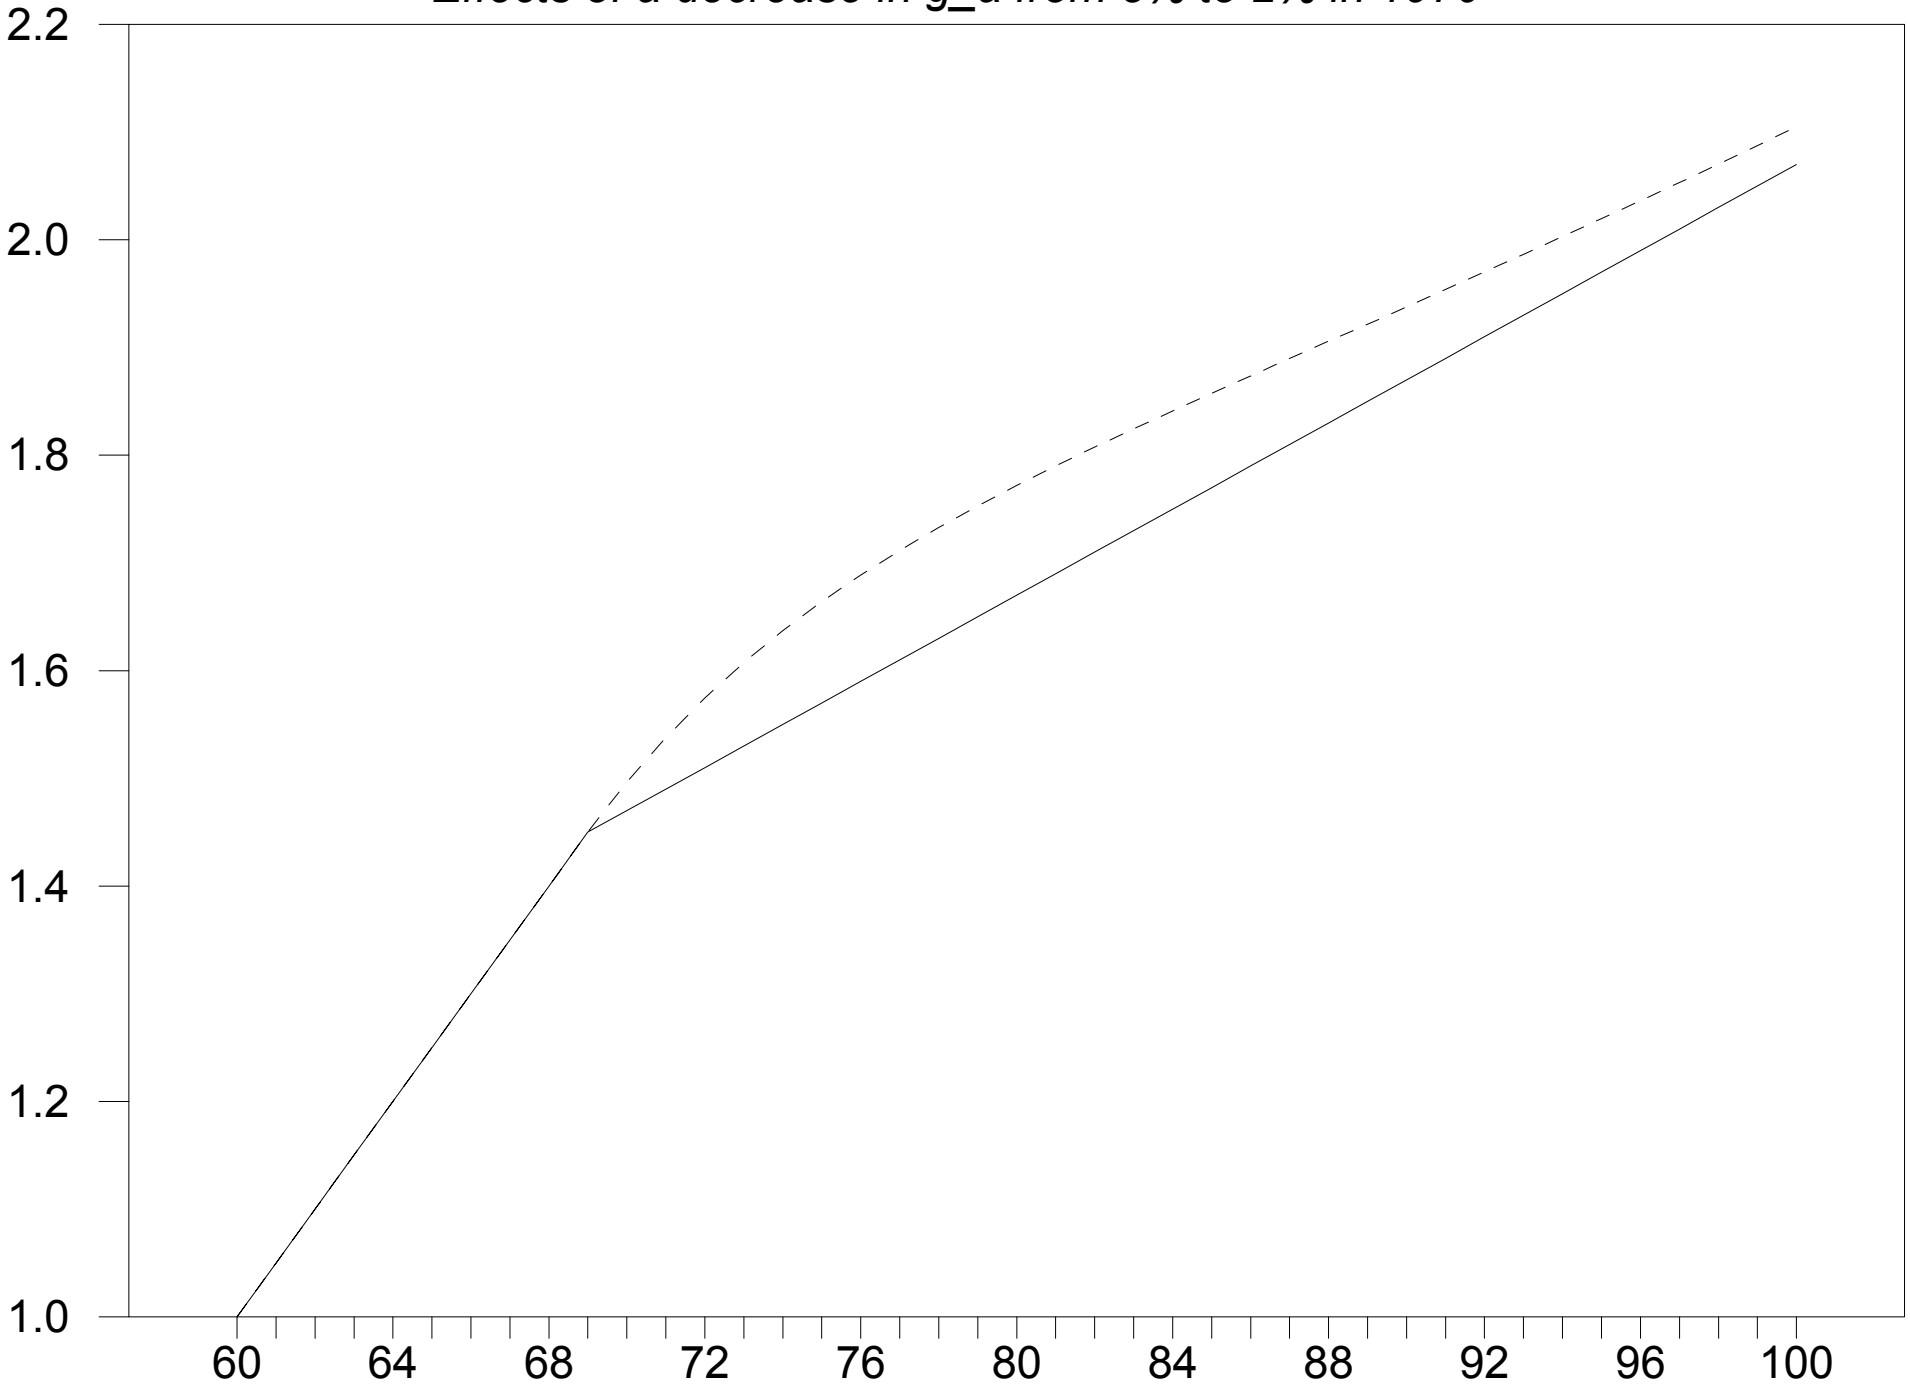

between (current and prospective) profit and user cost translate in capital accumulation or decumulation.

The parameters characterizing the cost of adjustment in the model imply an elasticity of investment with respect to the shadow price of capital equal to 1.0. Empirical evidence on the relation of investment to Tobin's Q yields lower elasticities. But, as discussed in that literature, these estimates are likely to be downward biased. More reliable, instrumental variable, approaches yield higher estimates. The study by Cummins et al. [1994] for example yields an elasticity around 1.0.

- The lower the costs of adjustment for factor proportions, the faster firms adjust the ratio of labor to capital to the wage.

  The parameters characterizing the cost of adjusting factor proportions imply a mean lag of adjustment of about 5 years. To get a sense of whether this is reasonable, think of the cost of adjustment as a short cut for a representation of the technology as putty–clay. In a world in which production were strictly putty-clay, only the newly installed capital stock, thus roughly 10% of the total capital stock each year (if we are thinking of equipment), would embody the new desired factor proportions. This would imply a mean lag of adjustment of 4.5 years.

Finally, the wage setting relation determines how unemployment in turn affects wages. The model assumes that an increase in the unemployment rate of one percentage point decreases the wage by 1%. For an unemployment rate of 10%, this corresponds to an elasticity of the wage with respect to unemployment equal to 0.1, roughly the number estimated by Blanchflower and Oswald [1994] in their estimation of a "wage curve" for a number of countries. My work with Katz [1997] has led me to conclude that the correct specification for the wage setting equation has richer dynamics than (1.4), and probably a higher long run elasticity than suggested by Blanchflower and Oswald. But I shall stick with the simple static specification here.

I solve the model under the assumption that nobody expected the slowdown in tfp growth before it happened, and that, when it happens, workers form their expectations as described above but firms have perfect foresight. In other words, firms understand that there has been a slowdown, understand that workers do not fully realize it but will eventually adjust their expectations over time. This is a strong assumption, but an assumption needs to be made. If firms are as confused as workers, then the effect on capital accumulation and on the labor–capital ratio will be smaller initially, but will last longer than those shown below.

The main results are shown in Figure 4 (When relevant, the variables are shown normalized by the level of technology, a.) The time unit is a year.

[Figure 4. The effects of a decrease in the rate of technological progress]

The wage increases for about 5 years, before eventually returning to its steady state value. The increase is small relative to the increase in (a^*/a): much of the effect shows up as higher unemployment, rather than as higher real wages. The evolution of the profit rate is the mirror image of the wage.

The higher wage leads firms to decrease the ratio of labor to capital over time. The ratio reaches a minimum after 9 years. And the lower profit rate leads to lower capital accumulation, with the trough taking place 16 years after the initial shock.

Lower capital and a lower labor–capital ratio both lead to a decrease in employment, and a corresponding increase in unemployment. The increase in the unemployment rate is largest after 10 years, equal to about 8%. The associated loss in output (relative to its balanced growth path value) 10 years out is equal to nearly 10%.

Of marginal interest here but of central interest for the next lecture is the behavior of the profit share. Despite the Cobb Douglas assumption, the profit share initially goes down, reflecting the fact that costs of adjusting factor proportions prevent an instantaneous adjustment of the labor–capital

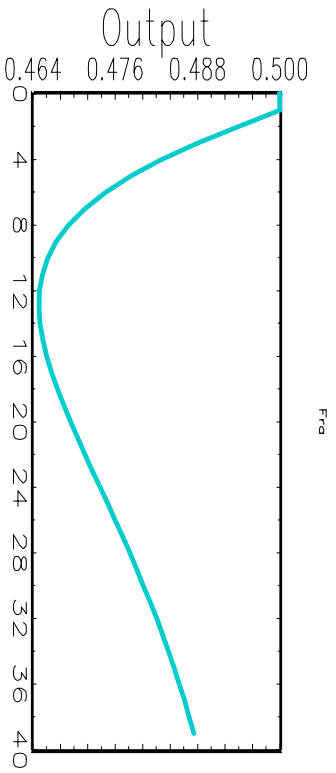
Output
0.464 0.476 0.488 0.500
0 4 8 12 16 20 24 28 32 36 40
Fra

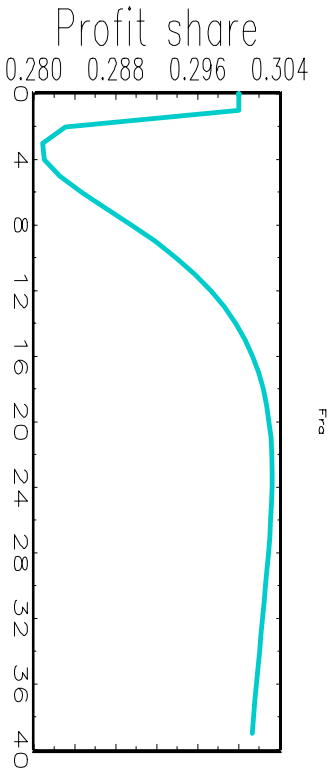
Profit share
0.280 0.288 0.296 0.304
0 4 8 12 16 20 24 28 32 36 40
Fra

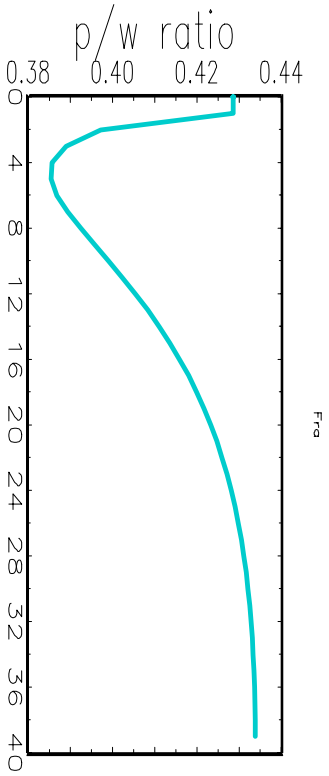
p/w ratio
0.38 0.40 0.42 0.44
0 4 8 12 16 20 24 28 32 36 40
Fra

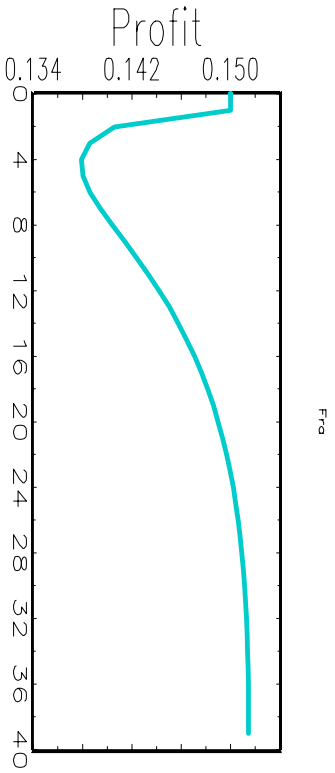
Profit
0.134 0.142 0.150
0 4 8 12 16 20 24 28 32 36 40
Fra

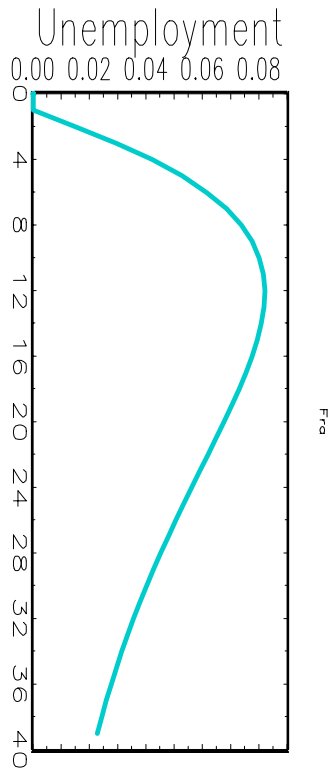
Unemployment
0.00 0.02 0.04 0.06 0.08
0 4 8 12 16 20 24 28 32 36 40
Fra

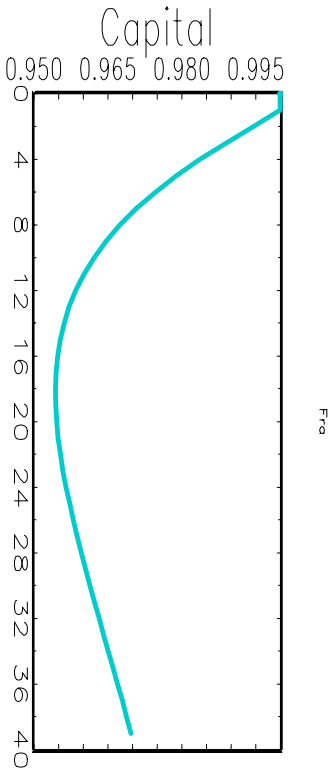
Capital
0.950 0.965 0.980 0.995
0 4 8 12 16 20 24 28 32 36 40
Fra

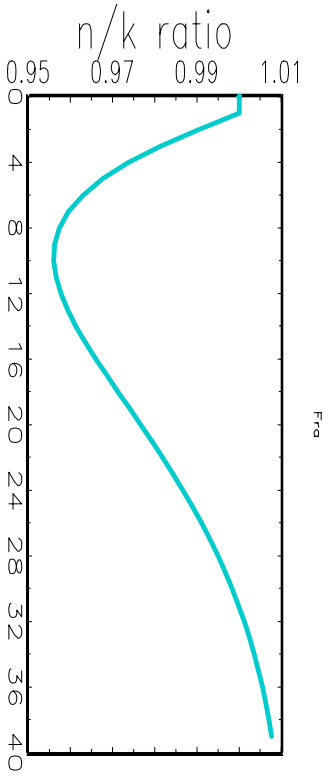
n/k ratio
0.95 0.97 0.99 1.01
0 4 8 12 16 20 24 28 32 36 40
Fra

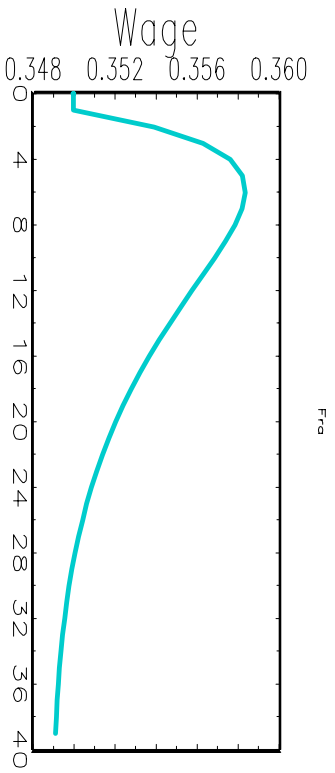
Wage
0.348 0.352 0.356 0.360
0 4 8 12 16 20 24 28 32 36 40
Fra

ratio. The share then recovers, overshooting slightly its steady state value before returning to it in the long run.

The limits of such a simulation are obvious, from the formalization of the adjustment of expectations, to the specification of the wage setting relation, to the choice of specific parameters. The purpose of the exercise was simply to see whether a decrease in the rate of technological progress could potentially generate a large and long increase in unemployment. I take the results to suggest that the answer is yes. According to this simulation, the slowdown in tfp growth in the mid–1970s can potentially explain much of the increase in European unemployment at least over the following 10 years, perhaps up to the mid or late 1980s. High unemployment has lasted much longer however. With this in mind, let me turn to a second exercise, the effects of higher interest rates on unemployment.

3 The effects of an increase in real interest rates.

The last 30 years have seen very large swings in real interest rates in Europe. For most European countries, the ex-ante real interest rate turned from positive in the 1950s and 1960s to sharply negative in the second half of the 1970s, and then to large and positive in the 1980s and much of the 1990s.[6] For some countries, the ups and downs of real interest rates have been quite dramatic: The real interest rate in Spain which had averaged about 2% in the 1960s decreased to –5% in the mid 1970s, and then back to 5% in the 1980s and the early 1990s.

This raises two issues. The first one is where these large movements came from. My own reading of the evidence points to monetary policy,

[6]This statement is based on a construction of ex-ante real rates as equal to the yield on medium-term government bonds minus realized inflation in the previous year. See Blanchard and Wolfers [2000] for details.

accomodating at first, much tighter later. The second is what effect these movements may have had on (equilibrium) unemployment. I shall start by leaving the first issue aside, and focusing on the second. I shall then return to the connection with monetary policy.

Consider an increase in the interest rate, followed by a return to its original value over time. (Why assume a temporary rather than a permanent change? Because, if we think of monetary policy as being behind the change, we want to keep the assumption of long–run neutrality, so the real interest rate eventually returns to its original level.)

Figure 5 shows how the economy reacts. The increase in the real interest rate implies an increase in the user cost. The fact that the profit rate is now below the user cost implies a decrease in capital accumulation, and thus a shift in the short run labor demand curve to the left, say from DD to DD' over time. The economy goes from A to B and unemployment increases. Then, as the interest rate returns to its original value, capital accumulation recovers, and the economy returns to its original equilibrium, A, over time.

[Figure 5. A temporary increase in interest rates]

The fact that an increase in the interest rate leads to capital decumulation and increased unemployment for some time is again not very surprising. The more interesting and again more difficult question is how large these effects might potentially be. And so, with the same caveats as for the decrease in technological progress earlier, let me embark on a simple simulation.

Assume that, until time $t = 0$, the interest rate r is equal to 5%. At time $t = 0$, it increases to 15%, and then returns to 5% over time according to:

$$r_t - 5\% = 0.9\ (r_{t-1} - 5\%)$$

Admittedly, a 10 percentage point change in the real interest rate is a

Figure 5. A temporary increase in interest rates

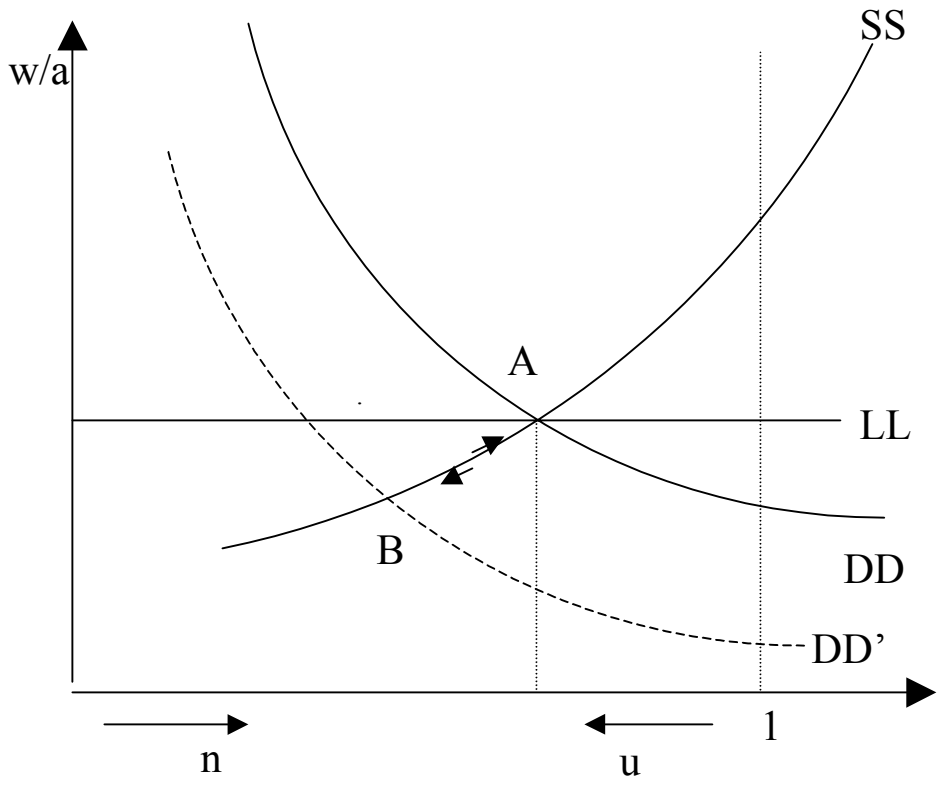

large one. But this is roughly what happened in Spain, first on the way down, from the 1960s to the 1970s, and, on the way up, from the 1970s to the 1980s.

The quantitative model is the same as before. The simulation assumes that the initial increase is unexpected, but, thereafter, both firms and workers anticipate the actual sequence of interest rates and thus the return to 5% over time. The results are shown in Figure 6.

[Figure 6. The effects of a temporary increase in interest rates.]

The increase in the user cost leads to capital decumulation. The effect is quite large, with a decrease of the capital stock (relative to its balanced growth path value) of close to 15% after 14 years. At a given ratio of labor to capital, capital decumulation leads to a parallel decrease in employment, and a corresponding increase in unemployment.

This increase in unemployment in turn leads to a decrease in wages, which leads to an increase in the ratio of labor to capital. Thus, there are two mechanisms affecting unemployment. First, the decrease in the capital stock, and, second, the increase in the ratio of labor to capital. The first must dominate the second, and the net result is an increase in unemployment.[7] Unemployment is higher by 5 percentage points after 8 years, before eventually returning to normal.

Again, marginal to our current interests but important for the next lecture is the behavior of the profit share. Because of costs of adjusting factor proportion, the ratio of labor to capital adjusts in response to the decrease

[7]The intuition, and the proof, is by contradiction: If the second effect dominated the first, unemployment would go down, leading to an increase in the wage. But if the wage increased, then the labor to capital ratio would decrease, not increase, a contradiction with the premise of the argument.

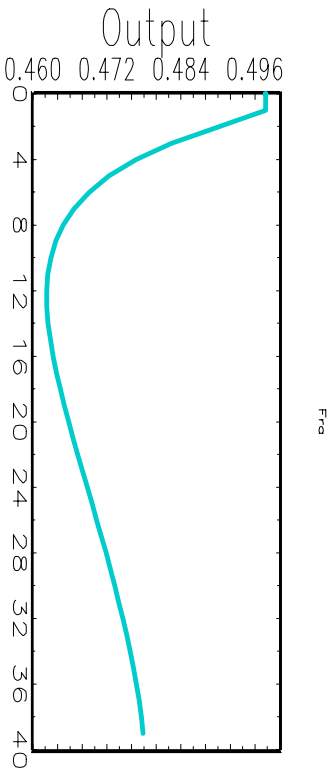
Output
0.460 0.472 0.484 0.496
0 4 8 12 16 20 24 28 32 36 40
Fra

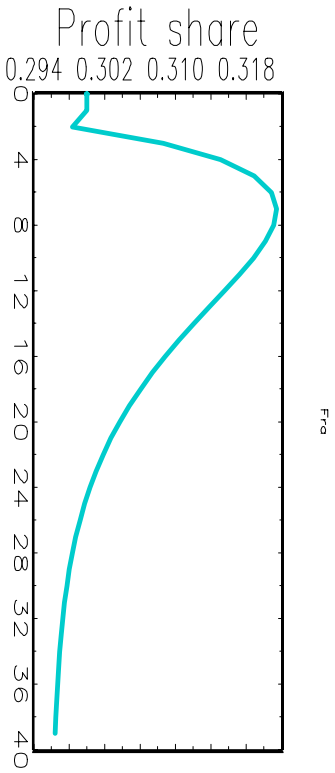
Profit share
0.294 0.302 0.310 0.318
0 4 8 12 16 20 24 28 32 36 40
Fra

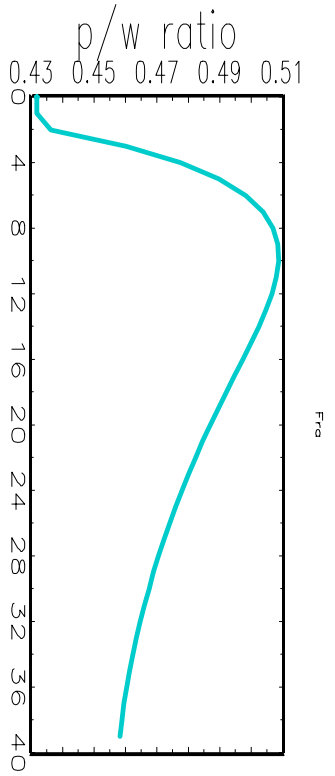
p/w ratio
0.43 0.45 0.47 0.49 0.51
0 4 8 12 16 20 24 28 32 36 40
Fra

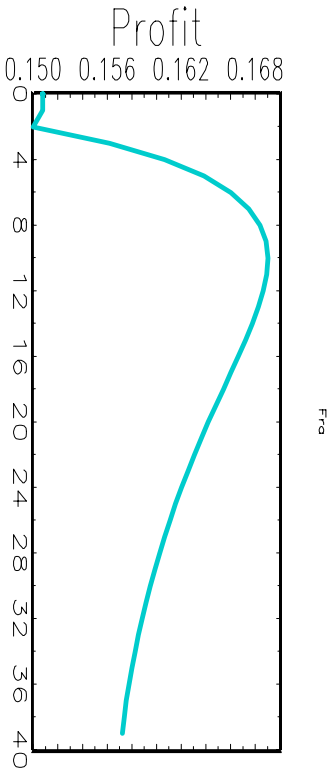
Profit
0.150 0.156 0.162 0.168
0 4 8 12 16 20 24 28 32 36 40
Fra

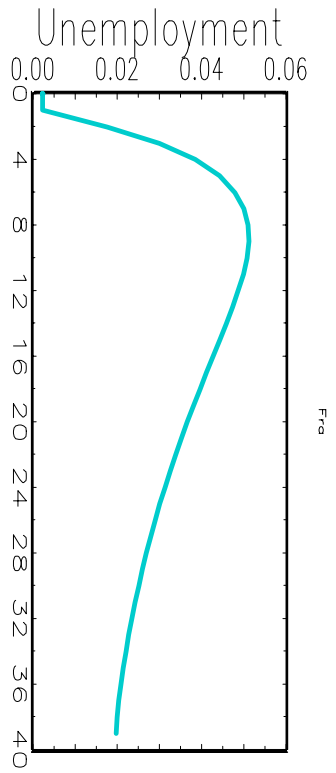
Unemployment
0.00 0.02 0.04 0.06
0 4 8 12 16 20 24 28 32 36 40
Fra

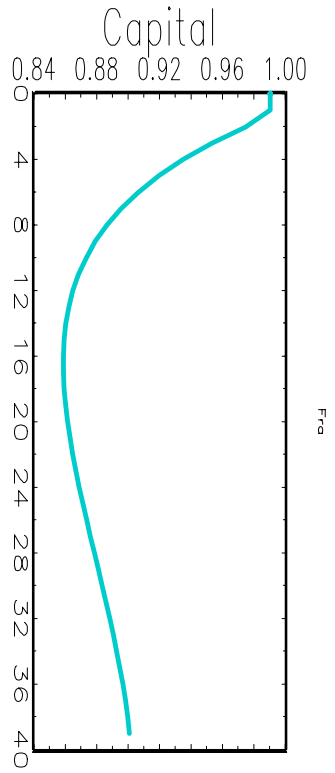
Capital
0.84 0.88 0.92 0.96 1.00
0 4 8 12 16 20 24 28 32 36 40
Fra

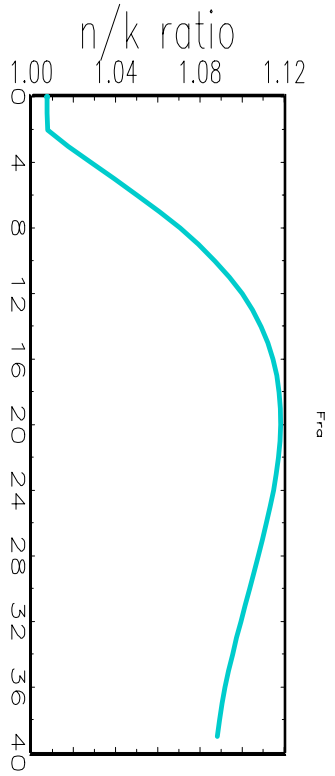
n/k ratio
1.00 1.04 1.08 1.12
0 4 8 12 16 20 24 28 32 36 40
Fra

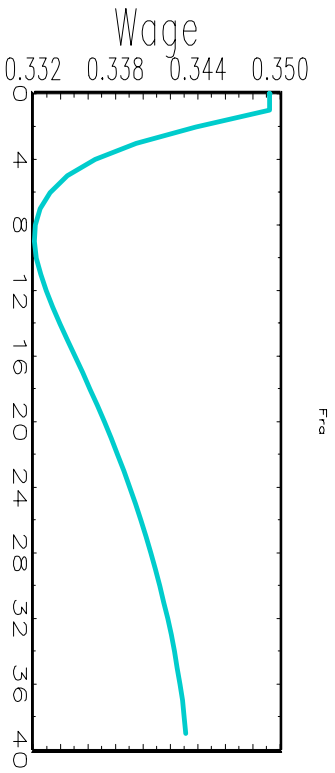
Wage
0.332 0.338 0.344 0.350
0 4 8 12 16 20 24 28 32 36 40
Fra

in the wage only over time, leading for some time to an increase in the profit share, followed by a return to its original value over time.

Going back to European unemployment: The argument developed in this section suggests that, some of the increase in unemployment which would have taken place in the 1970s was instead shifted to later, to the 1980s and maybe even to the early 1990s. Put another way, the evolution of interest rates helps explain why high equilibrium unemployment persisted into the 1980s, although it still leaves us short of a convincing explanation for why it persisted well into the 1990s in most countries.

Let me now go back to the other half of the question raised at the start of this section. Why did interest rates move the way they did? Many factors can lead to changes in real interest rates, from demographic changes, to changes in profitability, to changes in monetary policy.[8] My reading of the evidence points to a major role for monetary policy. When inflation increased in the 1970s, monetary policy was lax, letting nominal rates lag behind inflation. Then, starting with the Thatcher disinflation of 1979 and moving to the European continent in the 1980s, monetary policy was steadily tightened, nominal rates increased, leading eventually to the low inflation rates we observe today.

If this conclusion is right, this points to a serious shortcoming of the model I have used until now: The model is purely "real", with no nominal rigidities, and so no role for monetary policy in affecting real interest rates. The usual excuse for ignoring nominal rigidities is that the effects of money are sufficiently short–lived that, if the focus is, as it is here, on medium–term movements in unemployment, it may make sense to ignore them. If however movements in money can have sufficiently long–lasting effects on

[8]For two studies of the evolution of real interest rates up to the mid 1980s, see Blanchard and Summers [1984], and Barro and Sala-i-Martin [1990]. These would be fun to update.

real interest rates to affect capital accumulation and unemployment, then the usual excuse just does not hold. And if monetary policy affects not only actual but also equilibrium unemployment, we may need to revisit its role. If workers and firms were really confused about the slowdown in productivity in the 1970s, but the central banks were not (I am quite sure they were, so the argument will remain hypothetical), wasn't there an argument for lowering interest rates until aspirations adjusted and, in this way, lead to a more stable natural rate along the way? Or, to take another example, recent research has shown that inflation targeting makes good sense if the equilibrium rate of unemployment is given, unaffected by monetary policy. How is that analysis affected if the equilibrium rate itself responds to monetary policy?

4 A brief look at the panel data evidence.

How far does one actually get in trying to explain the evolutions of unemployment, both over time and across countries, based on the evolution of technological progress and real interest rates? This was one of the questions we took up in Blanchard and Wolfers [2000]. Let me summarize and sligthly extend our findings here.

We started by constructing, for each of 20 OECD countries, annual measures of labor augmenting technological progress and of real interest rates.[9] (We actually included a third variable, which we refered to as a shift in labor demand, or a shift in labor hoarding. I now prefer to think of this change as having its source in a change in institutions, and so I shall leave it out

[9]A semantic mea culpa: In that paper, we refered to the rate of labor–augmenting technological progress as the rate of total factor productivity (tfp) growth. The first is in fact equal to the second divided by the labor share. In these lectures, I am more careful about distinguishing between the two.

at this stage. But it will be the main focus on the next lecture.) We then constructed five–year averages of unemployment and each of the two shocks, for eight periods, starting in 1960-1964, and ending with 1995-1998.

We first ran a regression allowing the unemployment rate to depend on current, or current and lagged once, values of each of the two shocks, as well as a country effect:

$$u_{it} = c_i + a_1(L)ga_{it} + a_2(L)r_{it} + \epsilon_{it}$$

where i denotes the country, and t denotes the 5-year time period. The results of such a regression, allowing for just current values and for current values and one lag on one or both variables are given in Table 1. (The results are slightly different from those in the published paper because I include only two of the three shocks we included there. But, because the third shock is largely orthogonal to the first two, the general conclusions below are the same as in that paper). Table 1 yields two conclusions:

- First, the coefficients on the rate of technological progress and on the real interest rate are significant, both economically and statistically. In the specification allowing for current and lagged values of g_a variable, a 1% decrease in g_a leads to an increase in the unemployment rate of 0.5% in the first five years, 1.3% in the following five years. (One might want to impose the restriction that the effect of the tfp growth variable is zero in the long run. But the time series dimension of the panel is too short to learn much about these low–frequency dynamics.) An increase of 1% of the real interest rate leads to an increase in the unemployment rate of about 0.6 to 0.7%, with the effect happening mostly within the first five–year period.

- Second, the regression does a very poor job of fitting the cross–country dimension of the panel. This is shown in Figure 7, which plots the actual and fitted changes in unemployment, from 1970-74 to 1990-94.

Most of the fitted values are positive, reflecting decreases in the rate of technological progress and increases in the real interest rate in most countries over the period. But there is little cross-country correlation between fitted and actual values. To take an example (which has been a personal obsession for a number of years), Spain and Portugal have roughly the same fitted value, but very different actual increases in unemployment.

[Figure 7. Actual and predicted changes in unemployment. 1970-74 to 1990-94.]

Table 1. Technological progress, interest rates, and unemployment.

	Dependent variable: u_{it}				
on :	ga_{it}	ga_{it-1}	r_{it}	r_{it-1}	$\bar{R}^2$
	–0.49		0.71		0.56
	(–3.2)		(6.2)		
	–0.40	–0.71	0.50	0.09	0.61
	(–2.2)	(–4.7)	(4.5)	(0.8)	
		(with country specific coefficients d_i)			
	–0.57		0.76		0.56
	(–2.3)		(3.9)		

Period of estimation: 1960–64 to 1990–94. For the regression with country specific coefficients, the coefficient is the average of the country specific coefficients.

Our next step in Blanchard and Wolfers [2000] was to then explore a specification allowing the effects of the shocks to be a function of the labor

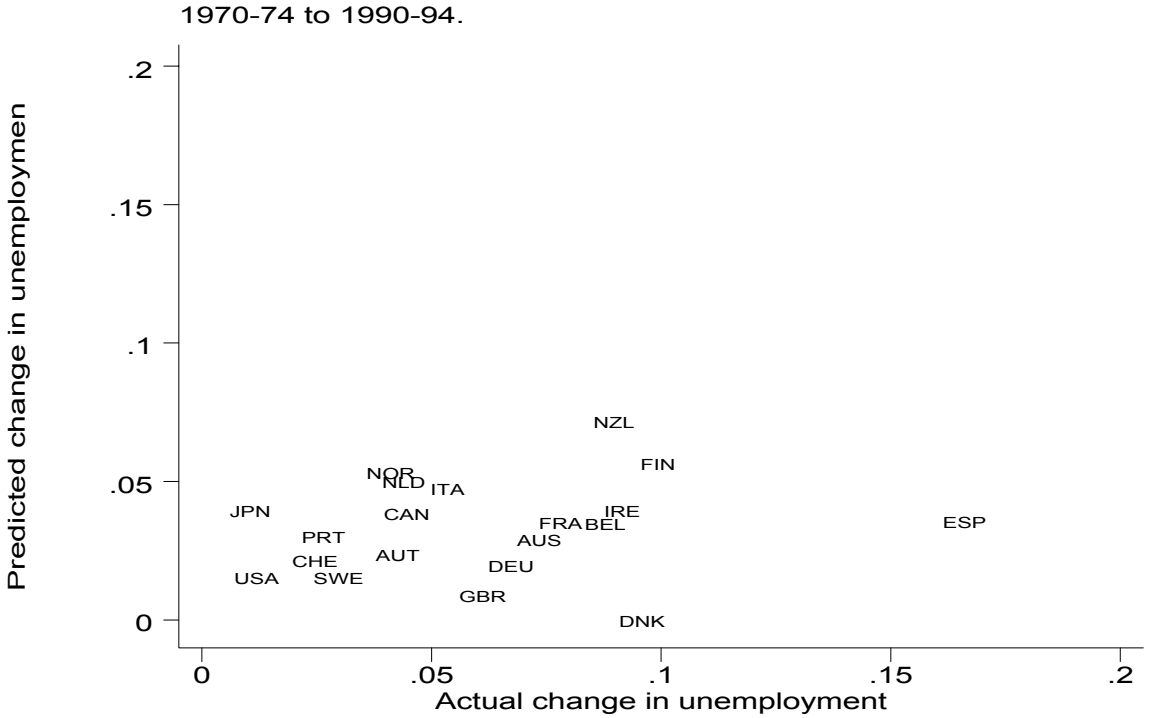

Figure 7. Actual and Predicted Change in unemployment

stata™

market institutions of each country. To avoid delving into institutions at this point, I shall present here the results from a slightly more agnostic approach, namely allowing the effects of the shocks to be country dependent.[10] The last line of Table 1 presents the results of a regression of the form:

$$u_{it} = c_i + d_i(a_1 ga_{it} + a_2 r_{it}) + \epsilon_{it}$$

where the effect of a given combination of shocks is now allowed to differ across countries. A normalization is required for the d_i's: The d_i's are normalized so their average across countries is equal to 1.

This regression yields three conclusions:

- There is again a significant effect of both tfp growth and real interest rates on unemployment.
- The d_i coefficients vary significantly across countries, from 2.3 for Spain to 0.4 for the U.S. This implies that a shock which leads to an increase in the unemployment rate of one percentage point in the average country leads to an increase of 2.3 percentage points in Spain, but only 0.4 percentage points in the U.S. Put another way, the equilibrium unemployment rate appears indeed much more stable in the U.S. than in Europe.
- The fit between predicted and actual values of the change in unemployment, shown in Figure 8, is much better than before. Most countries are close to the 45 degree line, although there are some clear exceptions, in particular Spain—where predicted unemployment increases by 10%, but actual unemployment increases by 15%. (The general

[10] A similar approach was taken by Phelps [1994], and more recently by Fitoussi et al. [2000].

fit becomes better, indeed becomes very good, when allowing for the third shock in the regression. See Blanchard and Wolfers [2000].)

[Figure 8. Actual and predicted changes in unemployment. 1970-74 to 1990-94. Country specific coefficients]

There are good reasons however to read these results with some skepticism. Allowing for the d_i to differ across countries is basically a way of allowing the regression to fit any increase in unemployment (if not the more subtle variations over time) for any given adverse shock. The reason I believe these results to contain some truth is the statistical relation of the estimated d_i's to a number of measures of labor market institutions. As we showed in Blanchard and Wolfers [2000], either forcing these coefficients to be a linear function of measures of institutions, or estimating them freely as done here and then regressing them on institutions, yields surprisingly plausible results, i.e. results which accord well with our priors of how institutions may affect the response of the economy to shocks. Although I shall not report and describe them, the same results hold here.

The conclusion I draw from this exercise is that one can indeed go a long way in explaining the evolution of European unemployment, at least up to the early 1990s, by the decrease in the rate of technological progress and the downs and ups of real interest rates.

Let me now turn to the more recent past, and see what light the model can shed on the recent declines in unemployment in countries such as Ireland and the Netherlands.

5 Ireland and the Netherlands.

Figure 9 shows the evolution of the unemployment rate in the Netherlands and in Ireland since 1970. In the Netherlands, the unemployment rate, which had reached 11% in 1983, has steadily decreased since, and is forecast to fall

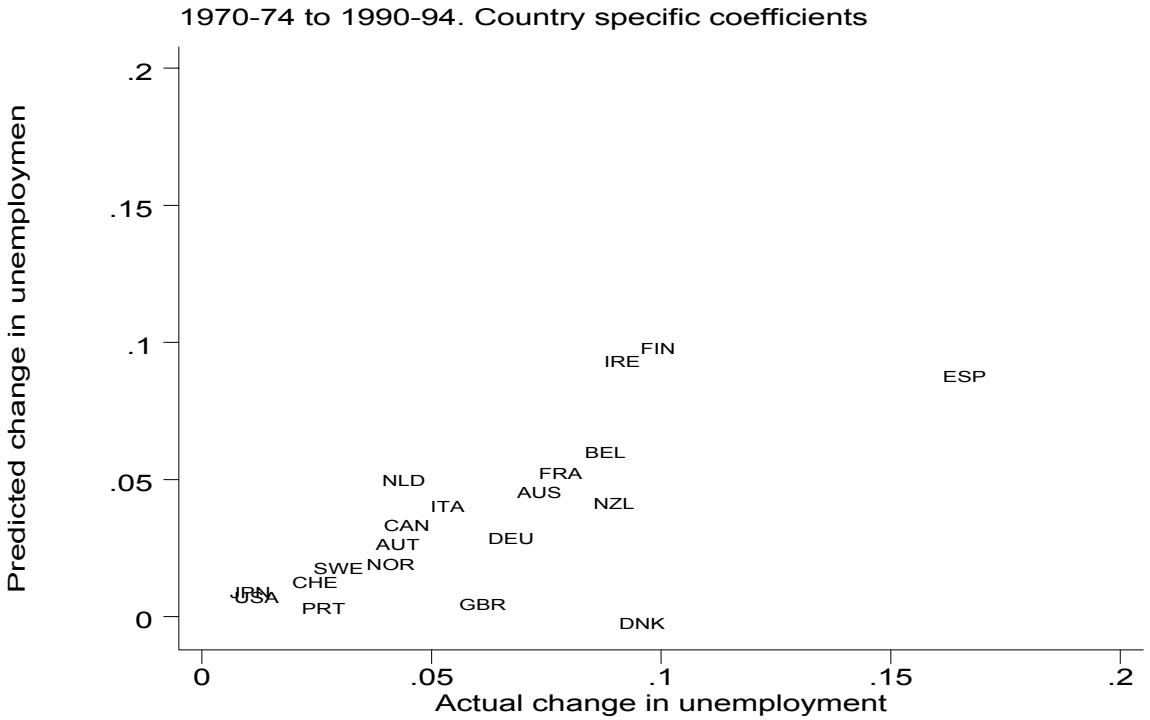

Figure 8. Actual and Predicted Change in unemployment

below 3% in 2000. In Ireland, the unemployment rate, which had reached 17% in 1986 and was still at 15% in 1993, has tumbled, down to 5% in 1999 and a forecast 3.6% in 2000.

[Figure 9. (a) Unemployment rate, Netherlands. (b) Unemployment rate, Ireland]

Such apparent "miracles" —as these two declines are often called— naturally raise doubts about their reality. But in neither case does the decline since the mid 1980s have much to do with statistical artifacts. It is often stated that there has been a large increase in part–time employment in the Netherlands, suggesting a shift in the distribution of work rather than a true increase in employment. Part–time employment has indeed increased, but this has been more than matched by an increase in the participation rate of women. It is also often mentioned that the number of officially disabled workers is suspiciously high in the Netherlands, hiding what should in fact be called unemployment. This is indeed true, but the proportion has declined since 1983, and so this cannot be the source of the decline in unemployment. In the case of Ireland, the evolution of unemployment actually understates the employment miracle: Employment growth has been larger than the decrease in unemployment, reflecting net in–migration, a sharp change from the net out–migration of the past.

Why have these two countries succeeded where others have not, or at least not yet? Many articles, and a few books, have already been written.[11] My goal here will be more limited, namely to look at the two countries

[11] Among them, for the Netherlands, the book by Visser and Hemerijck [1997], and the article by Nickell and van Ours [2000]. Broer et al. [2000] follows an econometric approach based on a model closely related to the model in this lecture. For Ireland, Fitz Gerald [2000] provides a very useful survey. Two good comparative studies (in French) are by Fitoussi and Passet [2000] and Freyssinet [2000].

Figure 9. Unemployment rate

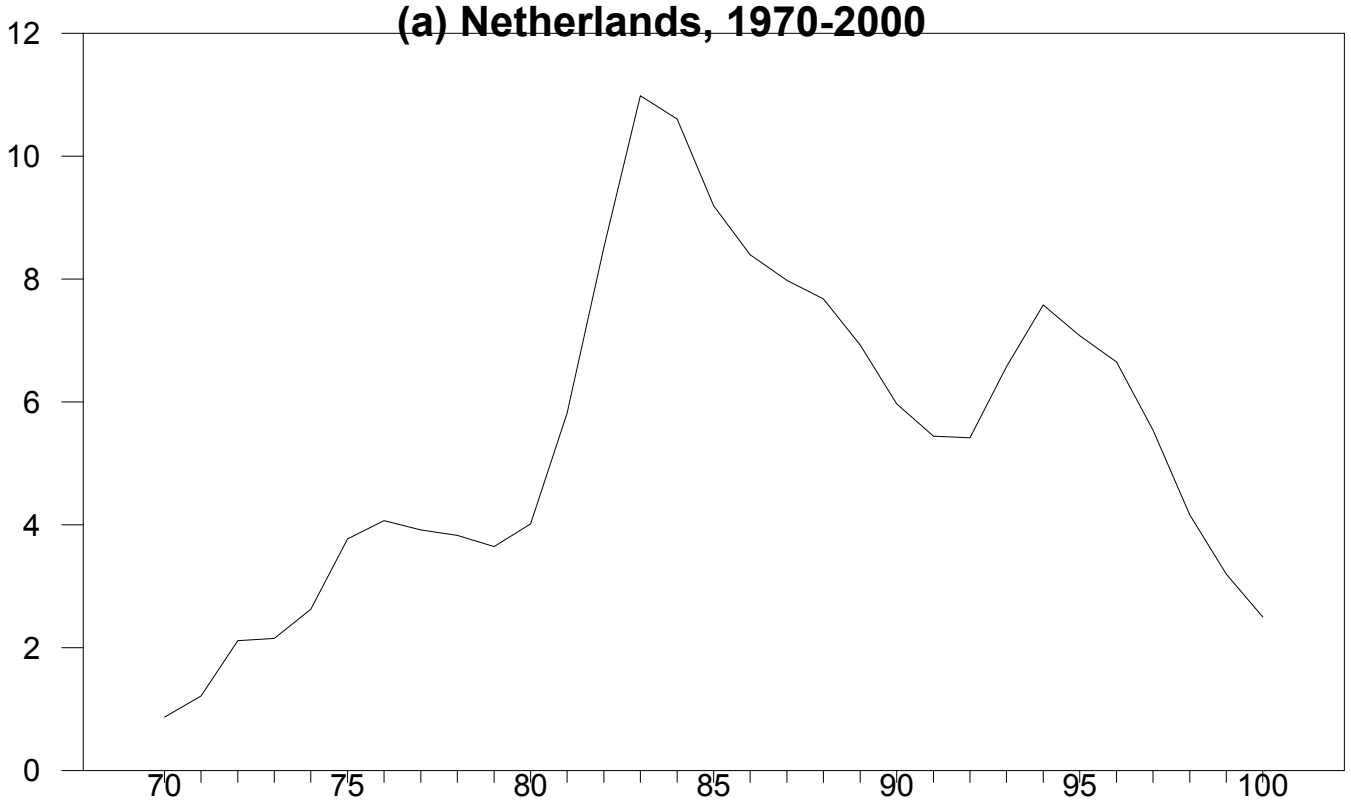

through the lens of the model, and see what this suggests. Let me state already the main conclusion: Wage moderation—in terms of the model, a substantial decrease in w/a—appears to be at the center of both evolutions. Its effects can clearly be seen on the labor–capital ratio, on capital accumulation, and on unemployment. The more difficult question is where this wage moderation has come from.

The model I sketched earlier suggests starting with the construction and examination of (w/a), the ratio of the real product wage to the level of technology—equivalently the wage in efficiency units. The model implies that, for the economy to remain on its balanced path with constant unemployment, the wage in efficiency units must remain constant, at a level consistent with the equality of the profit rate and the user cost. If (w/a) increases above this critical value, unemployment will steadily increase. If (w/a) decreases below this value, unemployment will steadily decrease.[12]

The first step is thus to construct a (Data sources and construction for both countries are given in the appendix.) This is done by integrating g_a over time, and taking the exponential of the resulting series.[13] The (logs of the) real wage w and of the level of technology a are plotted in Figure 10(a) for the Netherlands. The two log levels are normalized to 0 in 1970, based on the idea that unemployment was roughly constant, and the Dutch economy was roughly on its balanced path, at the time. The data here and in the figures below refer to the business sector only.

[12]If the economy does not exhibit Harrod neutral technological progress, the proposition still holds, but in slightly modified fashion. The wage associated with a constant profit rate and constant employment is still given by a constant value of w/a, where a is constructed in the same way as below. But, as there is no balanced growth path in this economy, this wage will typically not be associated with balanced growth.

[13]Recall that g_a itself is constructed as the Solow residual for each year divided by the labor share in that year.

The Netherlands

[Figure 10. Netherlands. (a) w and a (b) (an/k) versus (w/a)]

Figure 10(a) shows the clear change in the evolution of w relative to a which took place in the early 1980s. The wage, which had grown faster than a from 1970 to the early 1980s, has grown much more slowly since then. Since 1983, the wage has increased by 15% less than a, an instance of clear wage moderation.[14]

The model predicts that such wage moderation should first have led firms to shift from capital to labor. Figure 10(b) plots the ratio of employment in efficiency units to capital, (an/k), against the wage in efficiency units, (w/a). From equation (1.1), there should be a tight relation between the two (a static relation in the simple analytical model, a dynamic relation in a world in which firms adjust factor proportions only over time). And indeed there is. After a decrease in the 1970s, the ratio of labor to capital turned around in the early 1980s and has increased by more than 20% since 1982.

The model also predicts that the decrease in (w/a) should also have led to an increase in the profit rate, and, other things equal, to an increase in capital accumulation. In light of this prediction, the increase in the investment rate has been surprisingly small. Figure 11 shows why: Other things have not been equal.

[14]Because of data limitations, I have used bodies rather than total hours worked as the measure of employment in the construction of the Solow residual and of the real wage (the wage bill divided by employment). Because of the increase in part–time as well as shorter workweeks for full–time workers, total hours have grown less than bodies. If we were to use hours rather than bodies, the figure would show both slightly faster tfp growth and slightly faster real wage growth per hour than plotted in Figure 10(a). But the difference would remain the same.

Figure 10. Netherlands

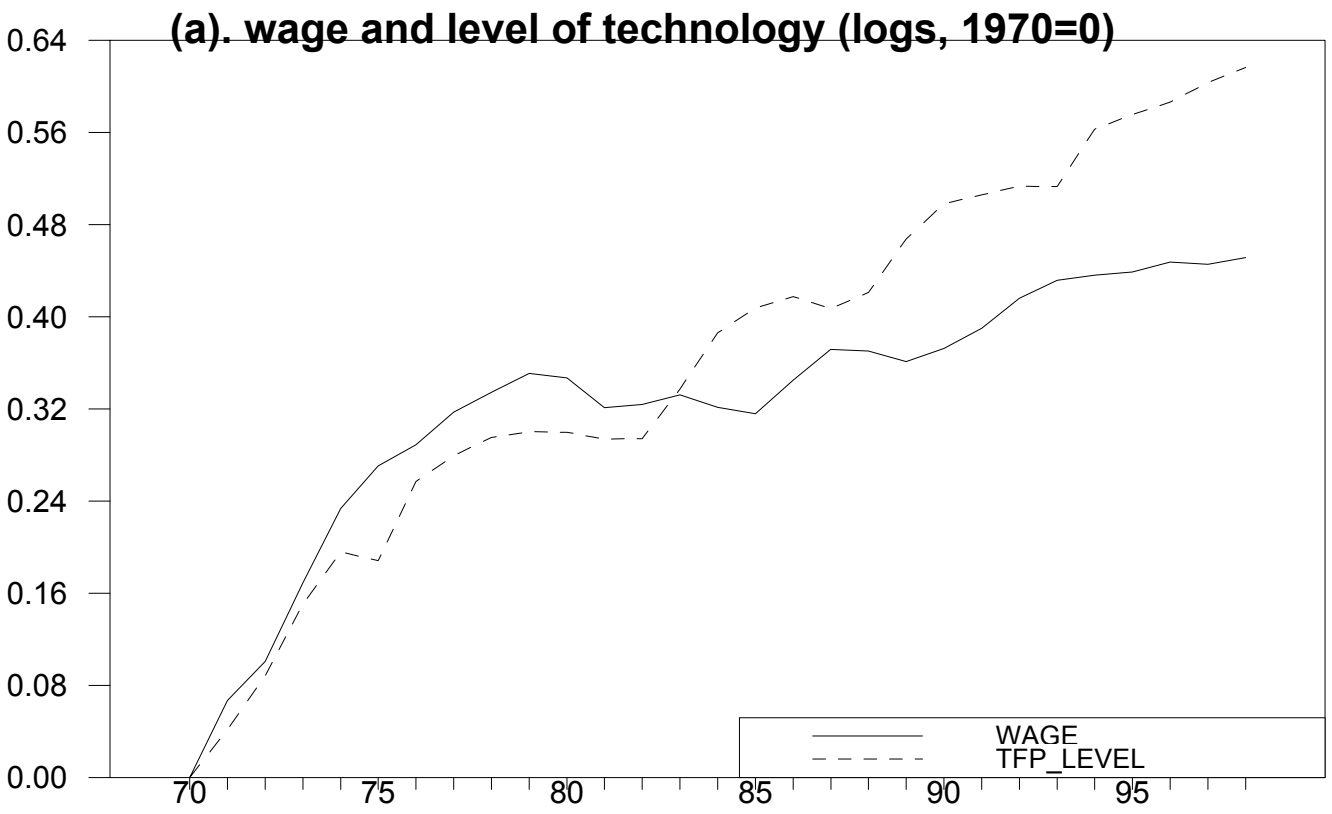

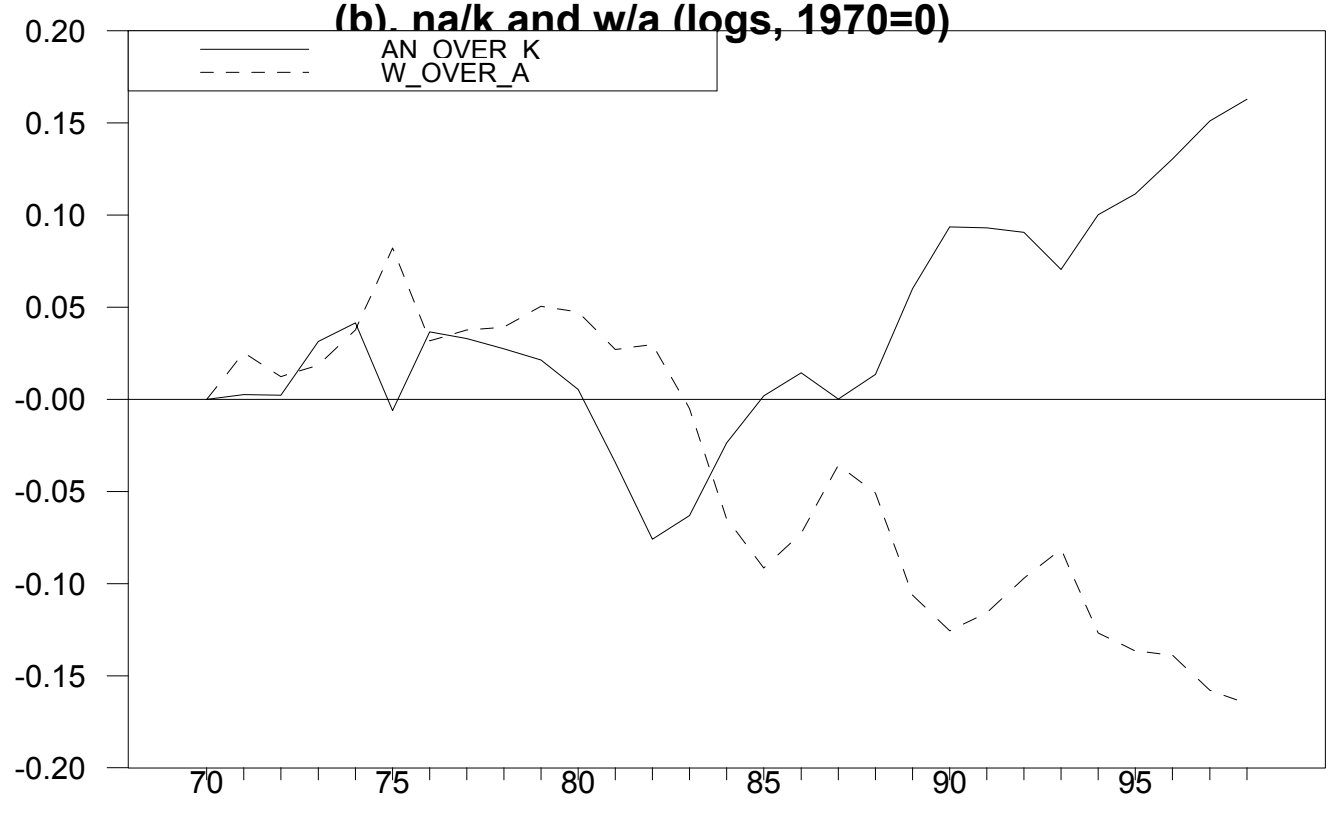

Figure 11(a) plots the profit rate (denoted rpi in the figure), i.e. the ratio of total profit to the capital stock in the business sector, and two measures of the user cost of capital. The first (denoted uc) ignores taxes, and constructs the user cost as the sum of the depreciation rate plus the real interest rate, itself constructed as the long nominal interest rate minus the average rate of inflation over the current and previous four years. The second (denoted uct) takes into account taxation along the lines of Jorgenson–Hall (a precise definition of the user cost in this case is given in the appendix.)

The figure yields two conclusions: First, the profit rate has indeed increased since the early 1980s. Second, this increase has been all but dwarfed by the increase in the user cost from the late 1970s to the early 1990s, an increase of more than 12 percentage points over 15 years. In other words, wage moderation has increased profits, but, until the early 1990s, this was more than offset by the increase in the user cost, due itself to a steady increase in real interest rates over the period.

This is what explains the mediocre performance of investment, at least until the early 1990s, as shown in Figure 10(b) which plots gross investment as a ratio to GDP. The investment rate, which had fallen in the 1970s has recovered only partly. (Another way to look at the evidence is to look at (k/a). (k/a) has remained roughly constant since the mid 1980s). The fact that the profit rate continues to increase, and the user cost to fall, in the second half of the 1990s suggests that an improvement might be coming. And, indeed, since 1997, gross investment for the economy as a whole (I do not have the investment data for the business sector yet) has grown at a rate 1.5% per year higher than GDP.

[Figure 11. Netherlands. (a) Profit rate and user costs. (b) Net investment rate]

To summarize, the proximate cause of the decrease in unemployment in the Netherlands appears to have been wage moderation. This has led

Figure 11. Netherlands

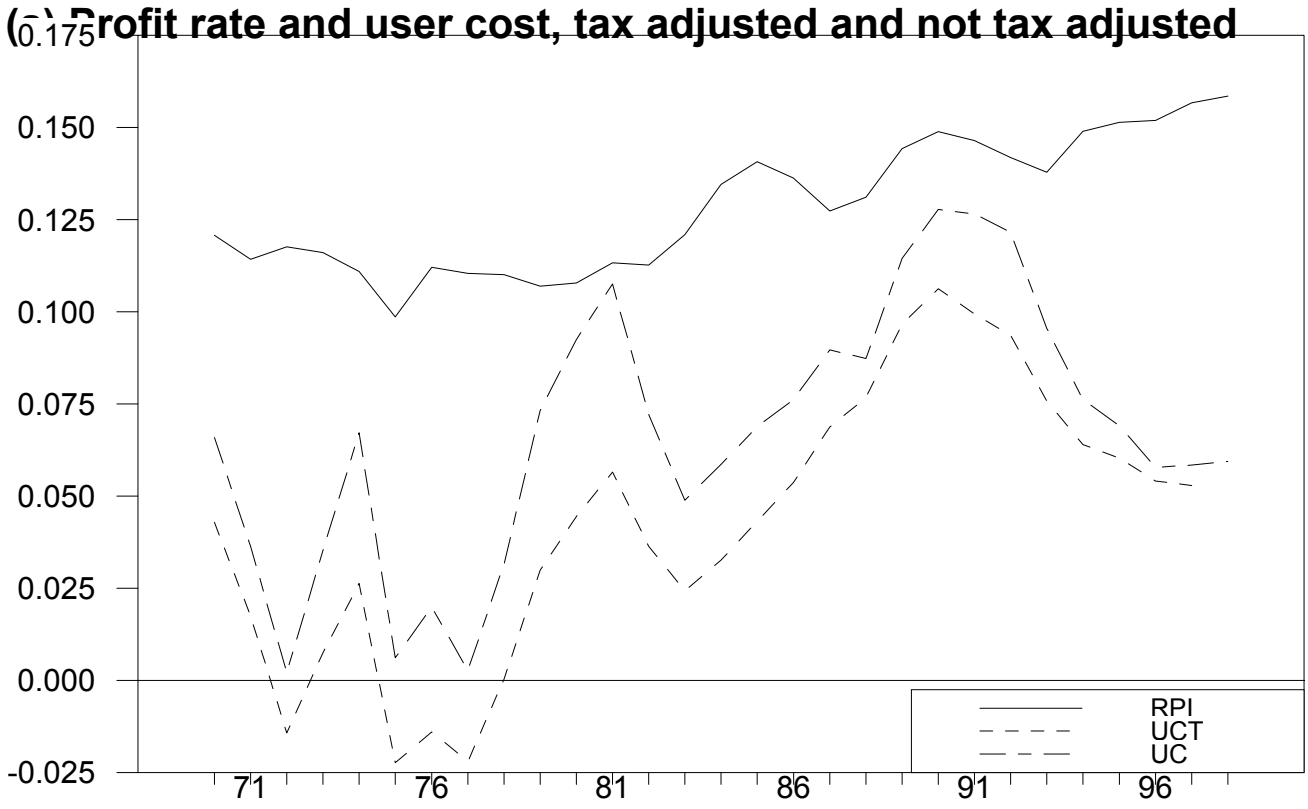

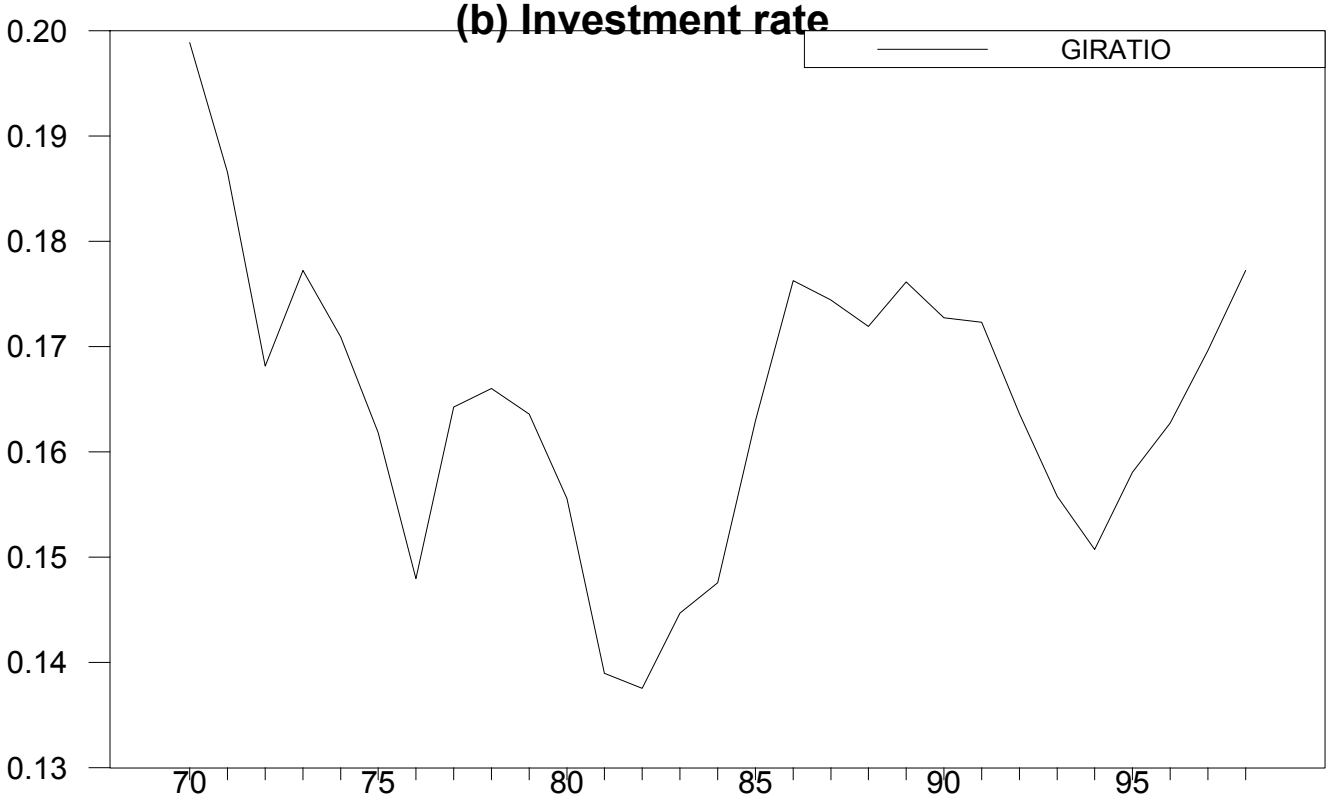

to an increase in the ratio of employment to capital. The effect on capital accumulation has been delayed by high interest rates; as interest rates have started to decline, the investment rate is now increasing. The increase in the ratio of employment to capital has led to an increase in employment, and a steady decrease in unemployment.

Ireland

Let me now carry out the same exercise for Ireland. Figure 12(a) shows the evolution of the real wage and the level of technology since 1970. The diverging evolution of the two since the early 1980s is again striking: Since 1983, w has increased by 40% less than a. (While the focus is on the decrease in unemployment, note another interesting aspect of the figure: There is no evidence of that excessive wage growth was at the root of the earlier increase in unemployment. The wage in efficiency units was roughly the same in the early 1980s as it was in 1970. This suggests either that the level of the wage was already too high in 1970, or/and that another factor was at work. As we shall see below, high real interest rates are indeed a plausible candidate.)

Figure 12(b) shows how, just as in the Netherlands but on a larger scale, this decrease in (w/a) has led to a turnaround of the ratio of labor to capital. Look at the scale on the left hand side: The ratio (an/k) has increased by more than 40% since 1983.[15]

[Figure 12. Ireland. (a) w and a (b) (an/k) versus (w/a)]

Turning to capital accumulation, Figure 13 shows that, even more than in the Netherlands, the performance of investment has been unimpressive. The explanation runs along the same lines as for the Netherlands. As shown in

[15]The disaggregated evidence suggests that this increase is more the result of a change in the mix of industries than a shift towards labor within industries (Bradley et al. [1993]); but the reasons and the results are basically the same.

Figure 12. Ireland

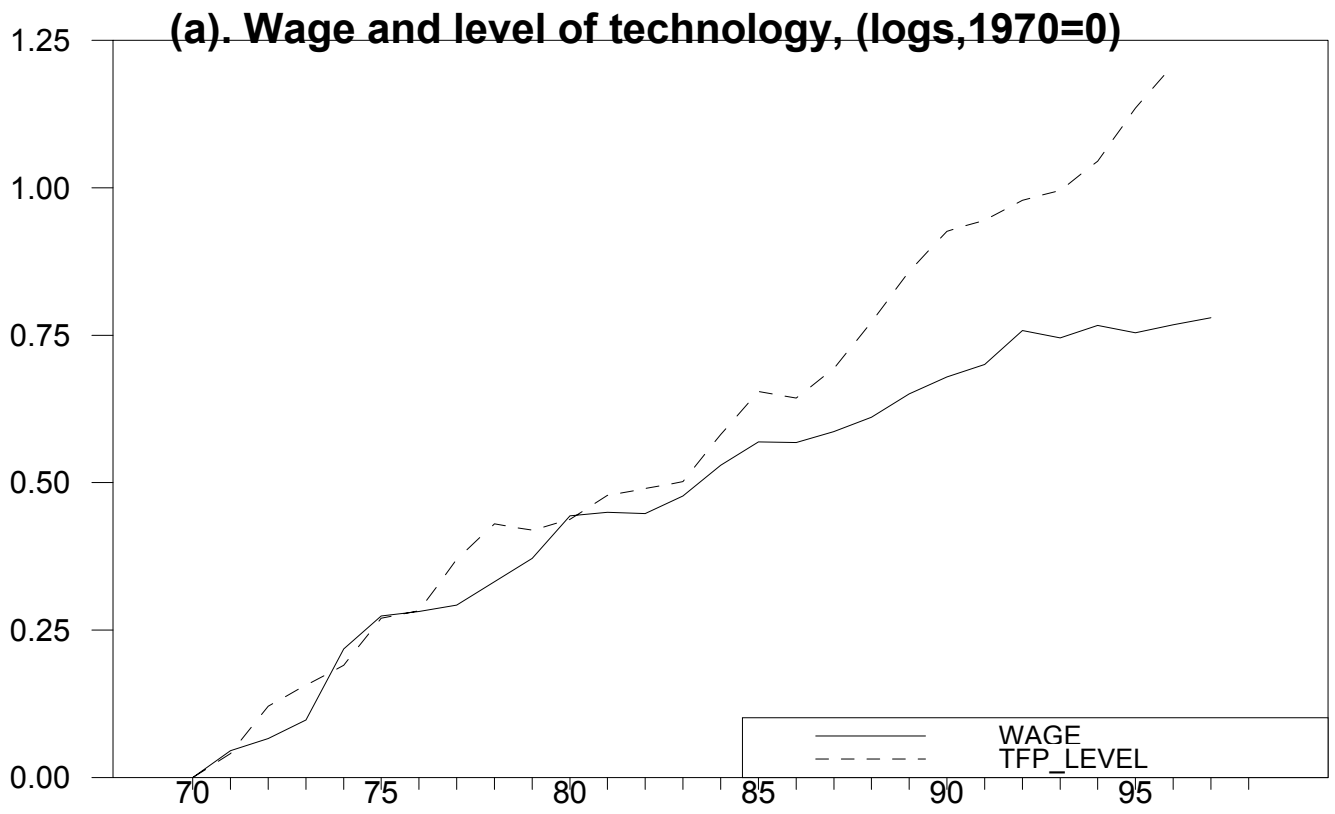

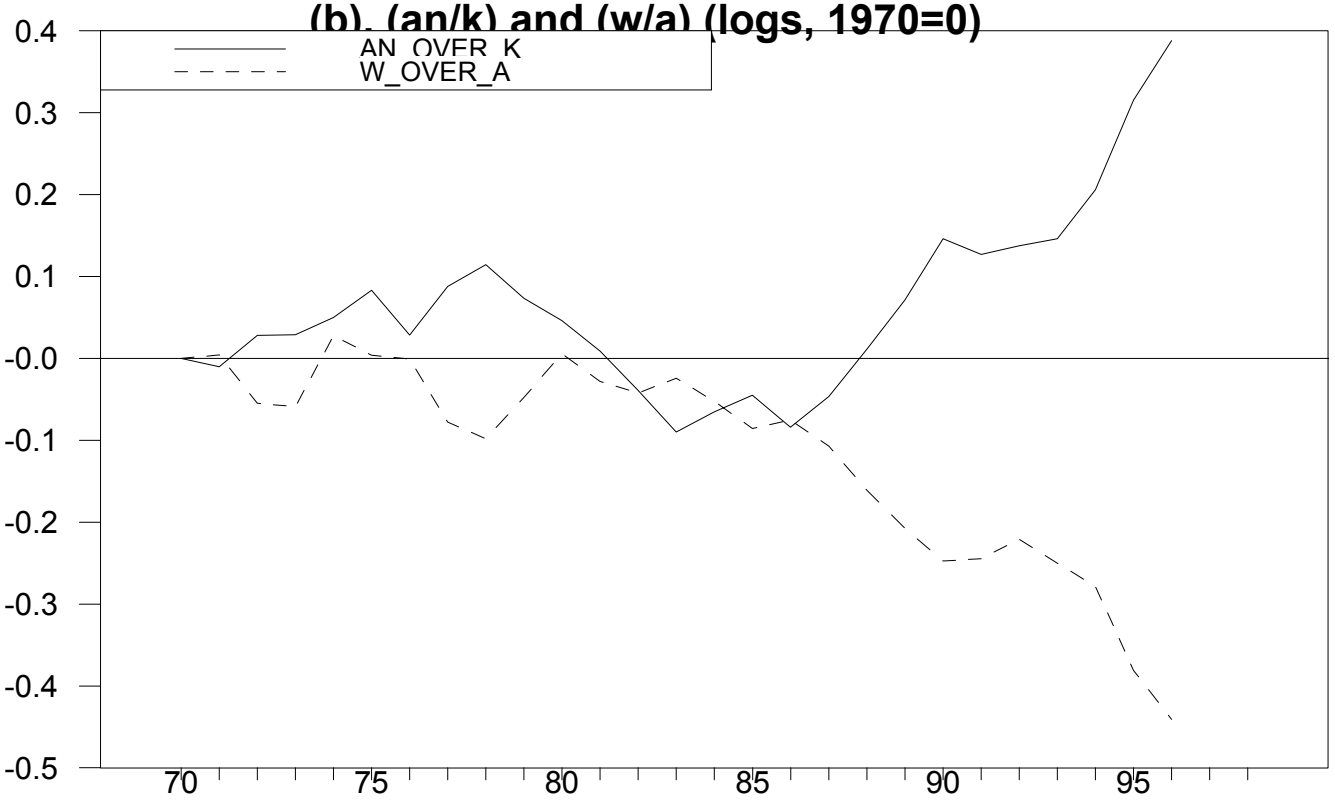

Figure 13(a), the profit rate has indeed increased steadily and considerably since 1983. But so has the user cost, from 1978 to at least the early 1990s. The reason behind the increase has been an increase in interest rates (rather than changes in taxation; despite the tax advantages given to foreign firms along the way, the evolution of tax–adjusted and non tax–adjusted series for the user cost are rather similar). And the increase in interest rates seems directly tracable to a change in monetary policy, i.e. to the entry of Ireland in the EMS in 1978.[16]

[Figure 13. Ireland. (a) Profit rate and user costs. (b) Net investment rate]

The recent further increases in the profit rate and decreases in the user cost suggest that Ireland should also see a large increase in investment. The data for business sector investment in Figure 13 end in 1996. But the evidence on the investment rate for the economy as a whole seems to bear this out. The gross investment rate has increased from 19% in 1996 to a forecast 23% in 2000.[17]

---

[16]This was indeed the perception at the time. See for example Dornbusch [1989].

[17]The very large increase in measured tfp growth (and by implication the rate of technological progress), the very large decrease in the wage in efficiency units, the very large increase in an/k, and the low measured rate of investment, all point to the possibility of measurement error in capital. If capital had in fact grown faster than reported, this would lead to a smaller increase in a, a smaller decrease in w/a, a smaller increase in an/k, and a more impressive investment performance. Looking into the issue, I found that the series for capital in the OECD business sector data base was indeed incorrect, based on an incorrect benchmark for capital at the start of the period, leading to very low growth of capital. I therefore use a corrected (lower) benchmark, based on Irish data, and the numbers presented here are based on this corrected benchmark (The details are available

Figure 13. Ireland

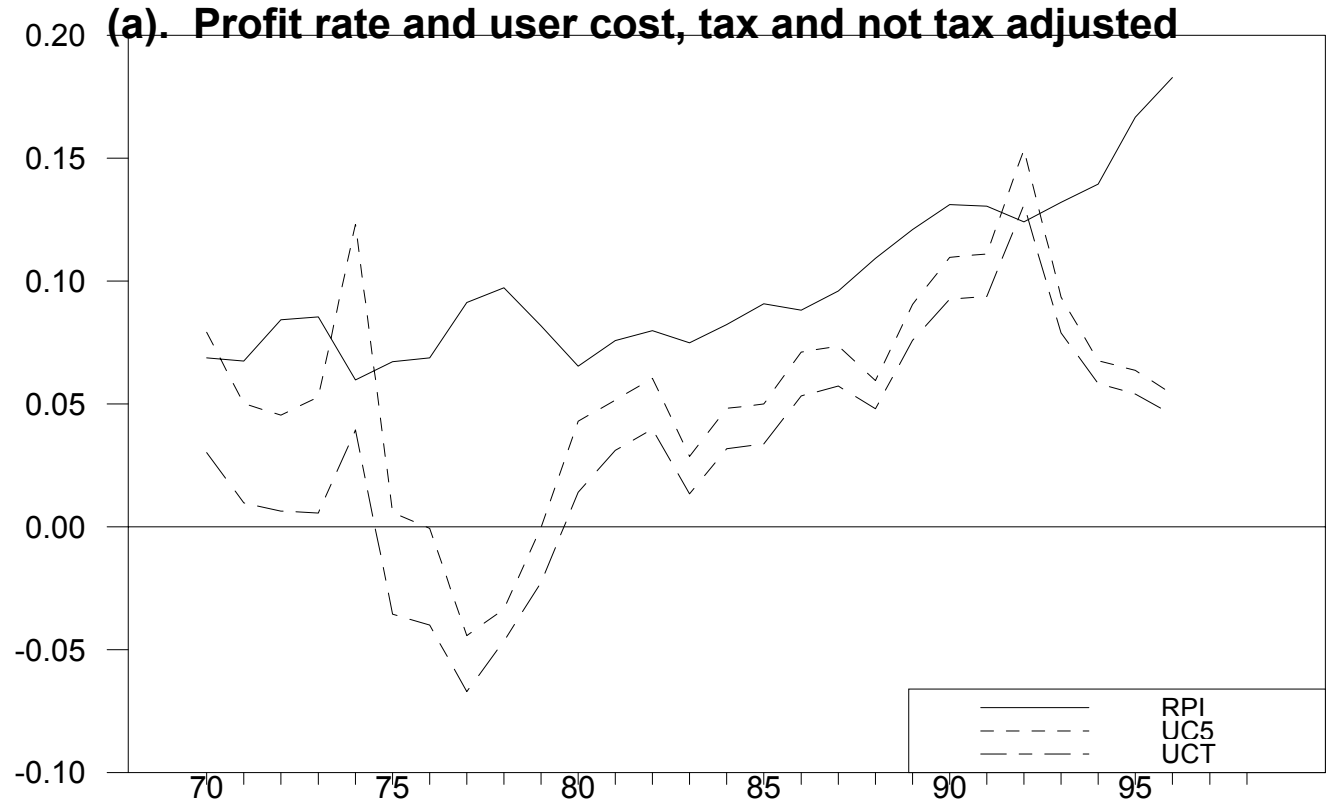

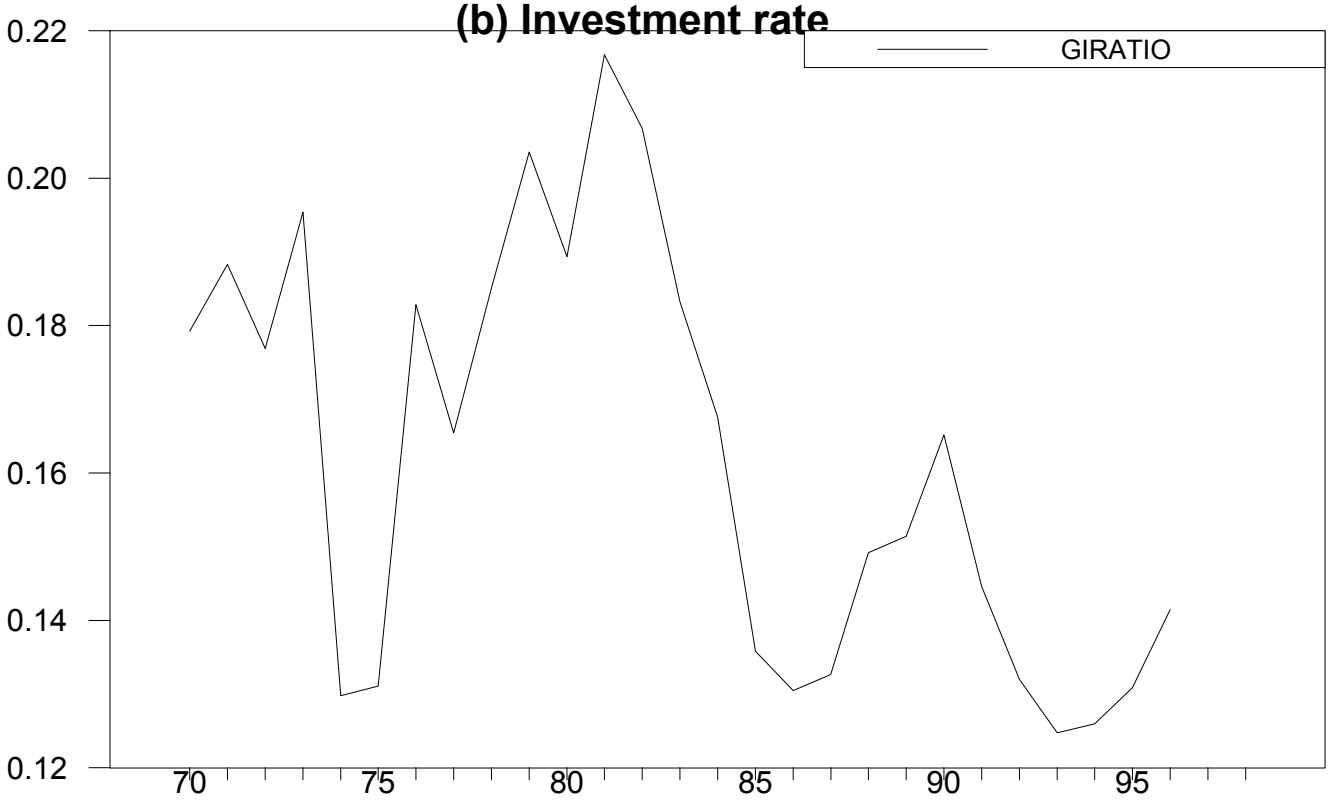

So the conclusions are roughly similar to those for the Netherlands. In both cases, wage moderation has led to an increase in employment relative to capital. The profit rate has increased, but until recently, so have user costs, leading to a limited increase in investment. The question this raises is an obvious one: What led to wage moderation? A full treatment would take me far beyond what I can do here. But I think the answer is somewhat different in each case.

Why the wage moderation?

The evolution of Ireland in the 1980s and 1990s seems to be the mirror image of our earlier analysis of the effects of the decrease in technological progress the 1970s. As can be seen from Figure 12(a), the rate of technological progress has increased in Ireland since the early 1980s. This is partly due to foreign investment, itself attracted by tax advantages, low labor costs, proximity to the markets of the European Union, and a skilled english–speaking labor force. But the effect has gone beyond foreign–owned firms, and tfp growth is high in most sectors of the economy.

The reason why this increase in technological progress has not led to a parallel increase in real wages is probably not misperceptions, but the high mobility between the U.K. and Irish labor markets. In–migration has limited the increase in wage growth in Ireland to a level close to that of the U.K., where the rate of technological progress has not increased, and therefore wage growth has been moderate. The recipe has worked well; but it can hardly be used by other countries, at least in this form.

The story is clearly different for the Netherlands, where, as can be seen from Figure 10(a), the rate of technological progress has not increased since

---

upon request). This implies higher capital growth than in the original OECD data. but even with the corrected benchmark however, capital growth is still surprisingly low.

the early 1980s.[18] What has happened instead has been a decline in wage growth.

Can one link this wage moderation back to changes in labor market institutions? The answer is probably not, at least not fully. While reforms, in particular a reform of the unemployment insurance system, have taken place, they do not appear to be substantial enough to account for the change in wage behavior since the early 1980s.[19] Most observers trace the start of wage moderation to the Wassenaar agreement, a tripartite agreement signed in 1982 between unions, business organizations, and the government. In that agreement, unions agreed to wage moderation (including the suspension of cost-of-living adjustments) in exchange for a number of measures ranging from a more generous financing of early retirements, and a movement towards a shorter workweek. Except for some wage pressure in the late 1980s, wage moderation has prevailed ever since.

This description leaves open however a number of deeper issues. Was the essential ingredient the presence of centralized bargaining?[20] The answer is not obvious. While centralized bargaining succeeded in 1982, it had been tried, without success, in previous years. At the same time, absent centralized bargaining and the complex package put in place by all sides

[18]If we were to measure employment by total hours rather than bodies, the measured rate of technological progress would be a bit higher, especially since 1983. From 1983 to 1997, annual hours per worker have declined from 1530 to 1365, a decline of 0.8% per year. I do not know the rate of decline from 1970 to 1983—which would be needed to know whether and how to modify the statement in the text.

[19]See the attempt at quantification by Nickell and van Ours [2000].

[20]The degree of centralization of bargaining is one of the measures of labor market institutions which, confirming earlier results by Nickell, we found in Blanchard and Wolfers [2000], leads to smaller effects of shocks on unemployment.

in 1982, wage moderation might have been much more difficult to achieve. Why did bargaining succeed in 1982, and not before? The rise in unemployment surely had something to do with it; that changes had to take place had become obvious to all.[21] And the increased intellectual and political acceptance of the notion that profitability was a key to a decrease in unemployment also played a role. Wage moderation went far beyond the usual effect of unemployment on wage demands. Whether one should categorize such a change in perceptions and the following wage moderation as a shock, in the same sense as the decrease in technological progress or an increase in interest rates, is unclear. But this is what has been at work.

6 Conclusions and extensions.

To summarize: The focus of this first lecture has been on shocks, and how they lead to changes in equilibrium unemployment. I have argued that the decrease in the rate of technological progress was at the root of the initial increase in European unemployment, that monetary policy shifted some of the increase from the 1970s to the 1980s, and that wage moderation is at the root of the recent unemployment miracles.

On this last point, an obvious question is how different evolutions have been in these "miracle" countries relative to the rest of Europe. The answer is that they have been indeed different. Wage moderation, as measured by the evolution of the wage in efficiency units, has much larger in Ireland and the Netherlands than in the rest of Europe. Still, there has been some wage moderation nearly everywhere, and this raises another issue: Why didn't this wage moderation lead to more of a decrease in European unemployment

[21] In the words of the Dutch Central Bank, "The government and the social partners reached the conclusion that "a limit was reached", and a change in policy and mentality was necessary". (De Neerlandsche Bank [1997])

than we saw in the 1990s. This apparent puzzle will be the topic of the next lecture. Before we get there however, let me, in this final section, take up three issues related to the theme of this lecture.[22]

Wages in efficiency units versus "real wage gaps"

In the early 1980s, a number of authors (among them Bruno and Sachs [1985]) suggested the use of "real wage gaps" as a diagnostic tool. Real wage gaps were defined as the deviation of the ratio of the wage to labor productivity $w/(y/n)$—equivalently the labor share wn/y—from some reference value, presumably a value where the ratio seemed consistent with balanced growth.

The idea was a simple and appealing one. If wages increased faster than productivity, labor costs increased, leading to a reduction in employment. But there were also obvious shortcomings. If, for example, the production function was Cobb Douglas, and there were no costs of adjusting factor proportions, the labor share would be constant, and by implication the real wage gap would always be equal to zero. In the presence of costs of adjustment, excessive wage growth would initially show up in a higher labor share, a positive wage gap. But, as firms adjusted factor proportions, the

[22]There is a fourth issue I would like to take up, an issue I have dubbed the "Modigliani puzzle". Franco Modigliani has shown that there is a surprisingly strong time–series relation between the net investment rate and the unemployment rate over the last 30 years both for the European Union as a whole as well as for many of the individual European countries (Modigliani [2000]. Modigliani's interpretation of this relation is in terms of aggregate demand, with investment demand determining demand, output, and by implication, unemployment. Given that the relation seems to hold at low frequency as well, I feel another explanation is needed. The question is whether it comes naturally out of the model sketched in this lecture. I do not have yet the data I need, and so this will wait until the next draft.

wage gap would disappear, while the problem was still there.

The measure of wages in efficiency units, w/a, avoids these shortcomings, and thus seems to be a more useful diagnostic device. It is not perfect, for at least three reasons.

- It is clear from the logic of the model presented at the start of the lecture that, in the long run, w/a is forced back to the value consistent with equality of the profit rate and the user cost. In the long run, what adjusts is unemployment. But, as the experience of the Netherlands and Ireland show, the dynamics here can be quite long, and the evolution of w/a can be quite visible for some time.

  Put another way, what one would like to have is a sense of changes in z in the wage relation $w/a = z\ h(u)$. What one observes is w/a, and in the long run, shifts in z have to be offset by movements in unemployment so $z\ h(u)$ remains constant. This suggests constructing a slighly more complex measure. Suppose that we are willing to assume that the wage relation takes the form:

  $$\log(w/a) = z - \beta u$$

  Then, if we were willing to assume a specific value for β we could construct a measure of z as $z \equiv \log(w/a) + \beta u$. This was the measure I constructed in Blanchard [1997]. While this extended measure is, I believe, very useful, it is also much more dependent on a particular specification of the wage relation.

- The second reason is empirical, and has to do with the need to use a series for capital. The constructed rate of tfp growth, and by implication, the rate of technological progress, depends on the rate of growth of the capital stock. And whether we measure this rate of growth at

all adequately is, for many countries, in doubt (see for example the discussion for Ireland earlier).

- The third reason is also empirical. Business cycle fluctuations may induce spurious movements in w/a. In the presence of nominal rigidities, a decrease in demand leads to a decrease in output, and, for reasons explored in the recent literature on the Solow residual, a decrease in measured tfp growth. Thus, a demand driven recession will lead to an increase in the measured w/a, potentially leading to the incorrect conclusion of a causal link from the wage in efficiency units to a decrease in employment.[23] The evolutions we have looked at in this lecture appear to happen at too low a frequency to be explained by business cycle fluctuations.

Nevertherless, despite these flaws, w/a seems like a useful measure to construct, and track over time.

Technological progress or labor productivity?

In thinking about the balanced growth path, it was a logical first step to define a wage setting relation in terms of a relation between the wage in efficiency units and the unemployment rate: Both are constant along the path. But this implies that workers think in terms of the underlying level of technological progress (or in terms of tfp growth, as the two are closely related), and this seems unlikely.

The question arises of how the conclusions I reached earlier would change if instead workers cared about the underlying level of labor productivity along the balanced growth path, if the wage relation took the form:

[23] Note the obvious parallel with the critique of RBCs as being based on a spurious relation between the Solow residual and output.

$$w = (\frac{y}{n})^{*} h(u)$$

And workers revised their perceptions of underlying productivity in the same way as before, but now with respect to labor productivity (y/n) rather than with respect to the level of technology a.

With respect to an underlying decrease in the rate of technological progress, the effect on unemployment would likely be stronger. The reason is that, as firms decreased their ratio of labor to capital in response to higher wages in efficiency units, labor productivity would increase relative to total factor productivity, thus slowing down the adjustment of expectations and leading to a larger and longer effect of the slowdown on unemployment.

In the limit, if there were no costs of adjusting factor proportions and, for simplicity, the production function was Cobb Douglas, firms would decrease the ratio of labor to capital so as to maintain labor productivity in line with wage growth. So if workers only looked at the evolution of labor productivity, they would never learn that there had been a decrease in the rate of technological progress. Given the lower profit rate, capital would decrease and the economy would implode until employment was equal to zero... The case is extreme, but the notion that, if workers care about labor productivity, attempts by firms to reduce costs by reducing the labor–capital ratio are likely to be partly self defeating, is very relevant.[24]

On the other hand, in response to a temporary increase in interest rates, the effect on unemployment would be weaker. The reason is that, as firms shifted away from labor, labor productivity would increase more slowly than

[24]The notion that, if wages respond quickly to labor productivity, things can go very wrong, is indeed an old theme. I remember an article by Martin Hellwig on Germany which makes this point, and which I have to trace. There is also a thoughtful discussion of these issues in Rowthorn [1998], some of which I agree with, some of which I do not.

tfp, moderating wage demands along the way.

To make progress here requires looking into the wage setting relation, and understanding how productivity affect wages, looking at what determines the reservation wage, at whether and how firms are able to reduce the hold–up problems involved in increasing the capital–labor ratio, at how workers and firms form expectations of the underlying level of productivity. I feel we are still a long way from a good empirical understanding of this essential dimension of wage setting.

Actual and equilibrium rates of unemployment

Robert Solow once remarked in jest that the best estimate of the natural rate of unemployment seemed to be the average of the previous three years of the actual rate. One can think of two explanations for this. One is that the actual rate affects the equilibrium rate, along the lines explored for example by hysteresis theories. The other, suggested by this lecture, is that the two rates do not depend on each other, but may move in similar ways in response to shocks.

Exploring the relation between the two rates with respect to shocks would require extending the model to allow for a distinction between actual and equilibrium rates, and thus for a role of aggregate demand in determining output. But one can guess where it would take us:

With respect to changes in monetary policy, both rates are likely to move in the same direction. The effect is likely to show up first in the actual rate, because of the effect of interest rates on demand, including investment demand. It is likely to show up in the natural rate over time, in response to lower capital accumulation. This may give the impression of the actual rate pulling the natural rate, where in fact what we observe are different dynamic effects on the two rates.[25]

The correlation is however likely to depend on the type of shock hitting

[25]Another implication of the fact that both rates are moving in the same direction

the economy. A wage explosion may for example increase aggregate demand and decrease actual unemployment, but increase the natural rate. This suggests one way of differentiating between causal or common dependence, which may be worth exploring in further work.

implies that disinflation may be more costly (in terms of unemployment) than derived in models which assume a constant natural rate—the same implication as under hysteresis, but through a different channel.

References

Barro, R. J. and Sala-i-Martin, X., 1990, World real interest rates, *NBER Macro Annual* pp. 15–74.

Blanchard, O., 1997, The medium run, *Brookings Papers on Economic Activity* **2**, 89–158.

Blanchard, O., 1998, Revisiting European unemployment: Unemployment, capital accumulation, and factor prices, Geary Lecture, Dublin 1998 (also NBER working paper 6566).

Blanchard, O. and Katz, L., 1997, What do we know and do not know about the natural rate of unemployment, *Journal of Economic Perspectives* 11(1).

Blanchard, O. and Summers, L., 1984, Perspectives on high world real interest rates, *Brookings Papers on Economic Activity* **2**, 273–324.

Blanchard, O. and Wolfers, J., 2000, Shocks and institutions and the rise of European unemployment. The aggregate evidence, *Economic Journal* **110(1)**, 1–33.

Blanchflower, D. and Oswald, A., 1994, *The wage curve*, MIT Press, Cambridge, MA and London.

Bradley, J., Fitz Gerald, J., and Kearney, I., 1993, Modelling supply in an open economy, *Economic modelling* **10(1)**, 11–21.

Broer, P., Draper, N., and Huizinga, F., 2000, The equilibrium rate of unemployment in the Netherlands, *De economist*, forthcoming.

Bruno, M. and Sachs, J., 1985, *The Economics of Worldwide Stagflation*, Basil Blackwell, Oxford.

Cummins, J., Hassett, K., and Hubbard, G., 1994, A reconsideration of investment behavior using tax reforms as natural experiments, *Brookings papers on Economic Activity* **2**, 1–74.

De Neerlandsche Bank, 1997, Zo kon het niet langer, *Jaaverslag* p. 24.

Dornbusch, R., 1989, Ireland's disinflation, *Economic Policy* pp. 174–209.

Fitoussi, J.-P., Jestaz, D., Phelps, E., and Zoega, G., 2000, Roots of the recent recoveries: Labor reforms or private sector forces?, *Brookings Papers on Economic Activity* 1, forthcoming.

Fitoussi, J.-P. and Passet, O., 2000, Reformes structurelles et politiques macroeconomiques: Les enseignements des "modeles" de pays, *Conseil d'Analyse Economique*.

Fitz Gerald, J., 2000, The story of Ireland's failure and belated success, mimeo, ESRI.

Freyssinet, J., 2000, La reduction du taux de chomage. Les enseignements des experiences Europeennes., *Conseil d'Analyse Economique*.

Modigliani, F., 2000, Europe's economic problems, Prepared for testimony before the Monetary Committee of the European Parliament.

Nickell, S. and van Ours, J., 2000, The netherlands and the united kingdom: a european unemployment miracle, *Economic Policy* 30.

Phelps, E., 1994, *Structural Slumps. The modern equilibrium theory of unemployment, interest, and assets*, Harvard University Press; Cambridge MA.

Rowthorn, R., 1998, Notes on unemployment and capital accumulation, (mimeo, Cambridge, England).

Visser, J. and Hemerijck, A., 1997, *A Dutch miracle. Job growth, welfare reform and corporatism in the Netherlands.*, Amsterdam University Press, Amsterdam.