Chapter 2:

More empirical results

2.1 Introduction

Chapter 1 examined whether the level of the minimum wage has a significant effect on the NAIRU. The econometric and other research used to answer this question has many further implications, some of which I explore in this chapter. The Table of Contents on page *viii* provides an outline of the chapter.

2.2 Other determinants of the NAIRU

2.2.1 The price wedge

As shown in Figure 1.1, the divergence between consumer prices and product prices, labeled *wedge*, has increased the NAIRU by about three quarters of a percentage point since the early 1980s. This section explores this effect in a little more detail.

Figure 2.1 below shows the wedge in level terms; that is, the ratio of the chain-weighted price index for personal consumption expenditure to that for output of the business sector, excluding housing and agriculture. The ratio is normalized to average 1 from 1947 to 1981, a period over which the two series rose by the same amount. Since the early 1980s consumption prices have risen, on average, by about half a percentage point a year faster than product prices. These broad patterns are reflected in the kinked trend in the ratio, also shown. In Chapter 1, the wedge is measured as the difference in this trend (in logarithms).

Figure 2.1: The price wedge

(ratio of consumption price to product price)

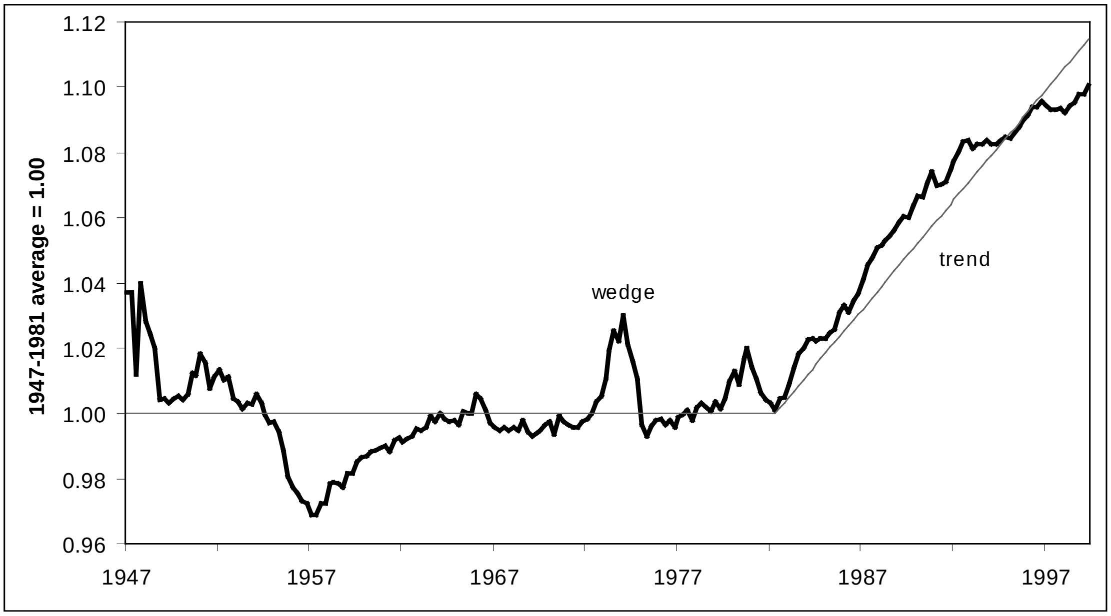

Movements in the wedge can be decomposed using National Income and Product Account identities. On the expenditure side, nominal GDP, P_yY, is the sum of consumption P_CC, investment P_II, government P_GG, and exports P_XX, less imports P_mM. That is:

$$P_yY = P_CC + P_II + P_GG + P_XX - P_MM \quad (1)^6$$

6 Equation numbering starts anew in each chapter.

The change in the price index for GDP can be approximated as the weighted sum of its components, with weights given by GDP shares:

$$\Delta\%P_Y = \Delta\%P_c s_c + \Delta\%P_I s_I + \Delta\%P_G s_G + \Delta\%P_X s_X - \Delta\%P_M s_M \qquad (2)$$

where $\Delta\%P_C$ represents the percentage change in consumption prices, s_c represents the share of consumption in GDP and so on.[7] This relationship can then be expressed in terms of relative prices by subtracting the change in the overall price level from each component. Denote the change in the relative price of consumption, $\Delta\%P_c - \Delta\%P_Y$ as $\Delta\%\tilde{P}_C$ and denote other relative prices similarly. After rearranging, the relative price of consumption can then be expressed as a weighted average of other relative prices:

$$\Delta\%\tilde{P}_C = -[\Delta\%\tilde{P}_I s_I + \Delta\%\tilde{P}_G s_G + \Delta\%\tilde{P}_X s_X - \Delta\%\tilde{P}_M s_M] / s_c \qquad (3)$$

GDP is also defined as the sum of sectoral outputs. These sectors are business excluding housing and agriculture (which sets the product price), farm, housing, government, and households and institutions. Denote the price of the output of each sector P_P, P_F, P_{HO}, P_{GO}, and P_{HI} respectively. Thus, a repetition of the above exercise

---

7 To see this, divide (1) throughout by Y, then difference, giving:

$$P_Y^1 - P_Y^0 = \left(\frac{P_C^1 C^1}{Y^1} - \frac{P_C^0 C^0}{Y^0}\right) + \left(\frac{P_I^1 I^1}{Y^1} - \frac{P_I^0 I^0}{Y^0}\right) + \left(\frac{P_G^1 G^1}{Y^1} - \frac{P_G^0 G^0}{Y^0}\right) + \left(\frac{P_X^1 X^1}{Y^1} - \frac{P_X^0 X^0}{Y^0}\right) - \left(\frac{P_M^1 M^1}{Y^1} - \frac{P_M^0 M^0}{Y^0}\right)$$

As price changes are large relative to quantity changes, this can be approximated as:

$$P_Y^1 - P_Y^0 \approx \left(P_C^1 - P_C^0\right)\frac{C^0}{Y^0} + \left(P_I^1 - P_I^0\right)\frac{I^0}{Y^0} + \left(P_G^1 - P_G^0\right)\frac{G^0}{Y^0} + \left(P_X^1 - P_X^0\right)\frac{X^0}{Y^0} - \left(P_M^1 - P_M^0\right)\frac{M^0}{Y^0}$$

The approximation is exact if the ratios of C, I, G, X and M to Y are constant.
Dividing throughout by P^0_Y and multiplying and dividing each term on the right hand side by its price gives

$$\frac{P_Y^1 - P_Y^0}{P_Y^0} \approx \frac{P_C^1 - P_C^0}{P_C^0}\frac{P_C^0 C^0}{P_Y^0 Y^0} + \frac{P_I^1 - P_I^0}{P_I^0}\frac{P_I^0 I^0}{P_Y^0 Y^0} + \frac{P_G^1 - P_G^0}{P_G^0}\frac{P_G^0 G^0}{P_Y^0 Y^0} + \frac{P_X^1 - P_X^0}{P_X^0}\frac{P_X^0 X^0}{P_Y^0 Y^0} - \frac{P_M^1 - P_M^0}{P_M^0}\frac{P_M^0 M^0}{P_Y^0 Y^0}$$

which is equivalent to (2).

can express relative product prices as a weighted average of the relative prices of other sectors:

$$\Delta\% \tilde{P}_P = -[\Delta\% \tilde{P}_F\ s_F + \Delta\% \tilde{P}_{HO}\ s_{HO} + \Delta\% \tilde{P}_{GO}\ s_{GO} + \Delta\% \tilde{P}_{HI}\ s_{HI}] / s_P \qquad (4)$$

subtracting (4) from (3) gives an equation for the wedge:

$$\begin{aligned} \Delta\%P_c - \Delta\%P_P &= \Delta\% \tilde{P}_C - \Delta\% \tilde{P}_P \\ &= [\Delta\% \tilde{P}_F\ s_F + \Delta\% \tilde{P}_{HO}\ s_{HO} + \Delta\% \tilde{P}_{GO}\ s_{GO} + \Delta\% \tilde{P}_{HI}\ s_{HI}] / s_P \\ &\quad - [\Delta\% \tilde{P}_I\, s_I + \Delta\% \tilde{P}_G\, s_G + \Delta\% \tilde{P}_X\, s_X - \Delta\% \tilde{P}_M\, s_M] / s_c \end{aligned} \qquad (5)$$

The panels in Figure 2.2 below show contributions to the wedge from the relative prices of major NIPA expenditure and output categories. The base period for each panel is from 1947 to 1981. The panels are drawn on a common scale, each series being weighted by its contribution to the level of the wedge. So, for example, the contribution of investment prices to the wedge is measured in Panel A as $-\Delta\% \tilde{P}_I\, s_I / s_c$ with the change being measured from the average level for 1947-1981.

Figure 2.2: Contributions to the price wedge

(in level terms, relative to 1947 – 1981 average)

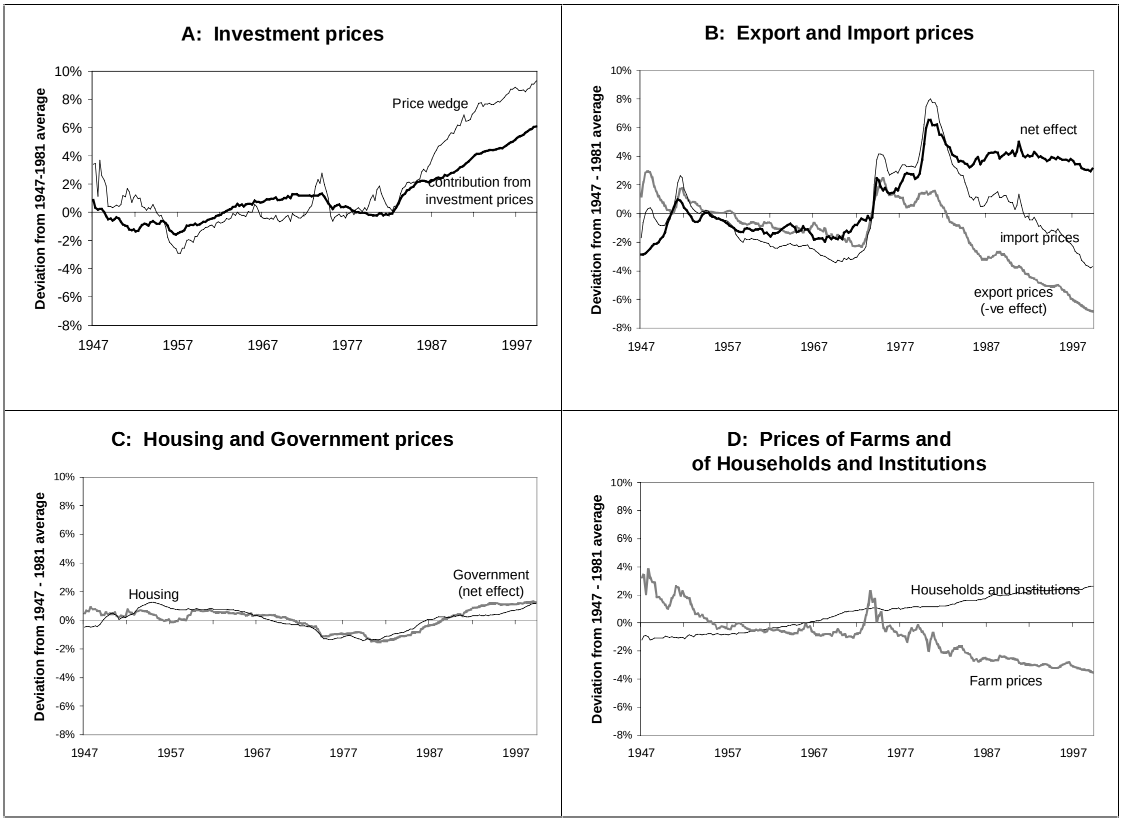

Panel A, which shows the contribution to the wedge from investment prices (with the level of the wedge also shown, for reference), is the most informative. The relative price of investment goods was relatively stable until the early 1980s, after which it fell steadily. Because investment goods are included in product prices but not consumption prices, this raises the wedge. Indeed, the relative price of investment goods seems to account for most of the longer term trends in wedge. This contribution, in turn, seems attributable to relatively rapid technological progress in the production of investment goods.

The ability to highlight explanations such as this is the main benefit of this accounting framework. The main limitation is the artificiality of separating effects like this from others. For example, because investment goods are mainly produced by the business sector (excluding housing and agriculture) a reduction in investment prices will raise the relative price of all other sectors. The accounting framework will then attribute part of the wedge to these other relative prices, even though it might be more appropriately described as an indirect effect of the relative price of investment goods. These interdependencies make interpreting other contributions to the wedge difficult. To some extent, common influences can be isolated by the presentation of net effects; however offsetting some contributions against others loses information and involves fairly arbitrary choices. Having made those caveats, the other major price series in the National Income and Product Accounts are shown in the remaining panels. (A small discrepancy, rising from the approximation involved in assuming constant GDP shares, is not shown).

Panel B shows contributions from export and import prices and their net effect, which can be interpreted as the contribution to the wedge from the external terms of trade. (The sign is reversed on the contribution of export prices). There is a trend deterioration in the terms of trade, which can be attributed to a trend decline in the relative price of agricultural output, much of which is exported. (Since 1967, when data became available, non-agricultural export prices have risen by about the same amount as import prices). This long term deterioration in the terms of trade accounts, in a formulaic way, for a large increase in the wedge over the last few decades.

However, this contribution is largely offset by a similar detraction from farm prices (Panel D). Deviations in the terms of trade from this trend partly reflect changes in the price of imported raw materials, particularly oil. These rose sharply in 1973 and again in 1979, causing pronounced short term increases in the wedge. Fluctuations in the terms of trade also arise from changes in the real exchange rate. A real depreciation raises import prices by more than export prices. This increases the wedge as the boost to consumption prices from imports is only partially offset by the boost to product prices from exports. Although trends in the real exchange rate are important in other countries, they do not seem to have a long term effect on the wedge in the United States.

Panel C shows contributions to the wedge from the output of the housing sector and from Government (measured as the sum of the effect from Government output and the effect from expenditures by Government, these two largely offsetting each other). Movements in both series are relatively small and show no strong trend. The correlation between the series seems to reflect the effects of unusual movements in other series on the benchmark. As might be expected, their relative prices move in the opposite direction to traded goods prices and investment good prices. This reinforces the contribution of investment good prices to the wedge (as noted above) but offsets that of the terms of trade.

Panel D shows contributions from output of the farm sector and that of households and institutions. The relative prices of both sectors show offsetting trends. Farm prices are highly variable in the short run, having large effects on the wedge in the late 1940s and mid 1970s, but these effects seem to be transient.

To conclude, movements in the relative price of investment goods can directly account for the timing and a large part of the size of the trend divergence between consumer and product prices. Sizeable deviations from this trend can be attributed to temporary increases in farm prices (in the late 1940s and mid 1970s) and to changes in oil prices (in the mid and late 1970s).

2.2.2 Unemployment benefits

There have been many studies of how unemployment benefits affect the incidence and duration of unemployment. As with the literature on the minimum wage, much of this is of questionable relevance. Flows in and out of unemployment – whether immediately attributable to unemployment insurance or not – need to be consistent with keeping unemployment near the NAIRU. It is only if benefits raise the NAIRU that they result in a permanent increase in unemployment.

My data set provides weak evidence that benefits raise the NAIRU. The statistical significance of the coefficient is not high. Consistent with this, the estimated effect varies with changes in specification or sample period. A more statistically robust effect can be found for unemployment benefits lagged three quarters, but this suggests overfitting. Averaging the series does not improve its performance.

That said, the estimated effect of unemployment benefits shown in Figure 1.1 or equation 8 of Chapter 1 is consistent with information from outside my data set. A wide range of theoretical interpretations imply that unemployment benefits will raise the NAIRU; Chapter 4 provides an example. And my estimate is similar to the effect estimated with international panel data by Nickell (1997 p64).

2.2.3 Productivity growth

Trend productivity growth has two effects on the NAIRU. A temporary effect arises because prices are assumed to respond to changes in the trend growth rate of productivity within a few quarters, while it takes about five years for the effect to flow through to wages. Thus the slowdown in productivity growth in the mid 1970s was initially fully reflected in higher growth in unit labor costs, boosting inflation and the NAIRU. Then, as wages begin to reflect the slower productivity, this effect fades. Even though this effect on the NAIRU is temporary, it seems important enough to warrant inclusion in Figure 1.1. The difficulty is that my estimates are imprecise and sensitive to arbitrary timing assumptions. Their presentation would be distracting rather than illuminating.

The more controversial and important issue is the permanent effect of productivity growth. Given the path of wages, faster productivity growth lowers unit labor costs. Given unemployment, this means less inflationary pressure and a lower NAIRU. However, this effect may be offset if nominal wages match increases in productivity. In the notation of Chapter 1, the offset is partial if $0 < \alpha < 1$ and exact if $\alpha = 1$, where α is the coefficient on productivity growth in the wage equation.

I estimate $\alpha = 0.52$ implying that, in the long run, about half a reduction in the trend rate of productivity growth is reflected in lower nominal wage growth. The other half increases the growth in unit labor costs and hence inflation and the NAIRU. This coefficient is not precisely determined, with a standard error of 0.20. I can reject the hypotheses that the effect is zero or one at conventional confidence levels. However I would not be able to reject hypotheses that the effect is near zero or near one. Nevertheless, a coefficient of about a half seems to be plausible and consistent with information outside my data set.

A positive effect of productivity growth on nominal wages is consistent with observations on individual wage adjustments. For example, a common standard for adjusting wages until the 1970s was "3 per cent plus COLA" (Mitchell, 1989, p193). And arbitrators explicitly refer to productivity changes in adjusting wages. It also appears to be consistent with bargaining theories of wage determination (see Pissarides, 1990 or Chapter 4).

A negative effect of productivity growth on the NAIRU is more controversial. Stiglitz (1997 p7), for example, writes "Neither the level or rate of change of productivity has any long run effect on the unemployment rate – witness the fact that unemployment has been about the same over the course of a century of massive productivity growth and large shifts in its trend growth rate".

My estimates of the NAIRU in the USA appear to be inconsistent with Stiglitz's observation. However, with only one assumed observation of a change in trend productivity, the results are not compelling. More persuasive evidence may lie in the experience of other industrialized countries – most of which suffered a reduction in trend productivity growth at the same time as that in the USA, followed by substantial increases in the NAIRU. Increases in the NAIRU varied from country to country, suggesting that institutional determinants are important. However there also appears to be a strong common element. (The positive vertical intercept of Figure 1.2). A slowdown in the rate of technological progress is an obvious candidate for explaining this worldwide shift. An alternative explanation, increases in energy prices, lost its force after energy prices collapsed in the mid 1980s.

In the long run, aggregate real wages appear to be proportional to productivity. However, because prices adjust so as to maintain this nexus (see Appendix 1 of Chapter 1), it is not necessary that wages do. At an industry level – computer hardware being a vivid example – the rapid adjustment of prices to productivity precludes much adjustment of wages.

There are many interesting questions that could be asked concerning the effect of productivity on the NAIRU. However, the lack of variation in the United States time series means that other data sets are likely to be more informative about these. In the absence of clear guidance regarding timing, smoothing or controlling for other inputs, I use a simple, widely used measure that fits the data. Specifically, I measure trend productivity as a linear trend fitted to the logarithm of output per hour in the non-farm business sector, with a kink at 1973:1. This is the same measure that successfully enters price equations (for example, that in Appendix 1 of Chapter 1). Differencing this series essentially creates a step variable (or a "post 1973 dummy"). To smooth the jump at the time of the kink, I take a 5 year back-average of the difference. This is consistent with the view that wage setting standards are slow to adjust to shifts in productivity growth and fits the data better. Alternative measures, such as a moving back-average of actual productivity, perform similarly in the wage equation, though less well in the price equation.

2.2.4 Changes in payroll taxes

I measure payroll taxes as the change in the statutory rate of employers' contributions to Old-Age, Sickness, Disability and Hospital Insurance (OASDHI). This has risen from 1 per cent in 1949 to 7.65 per cent currently. In principle, I would also like to include mandatory contributions to unemployment insurance (1 per cent of wages in 1996) and workers compensation, however doing so is not simple. (While some data on these charges are available, to express them as a rate involves using the dependent variable in the denominator, giving rise to simultaneity bias). It would also be desirable, in principle, to take explicit account of changes in OASDHI exemptions, though given that these are typically indexed, such an adjustment might make little difference.

My estimates imply that a percentage point increase in the employer's OASDHI rate raises average compensation (including OASDHI) by 0.88%. Because employer contributions approximately equal employee contributions, this implies that the immediate incidence of these taxes is similar to the legal impact. Employees pay the employee contribution while employers pay almost all the employer contribution. Possible reasons for the coefficient not equaling unity include exemptions, non-compliance, shifting of the burden from employers onto employees and specification bias.

To offset the impact of this payroll tax increase would require an extra percentage point of unemployment for six quarters. Averaged over the sample period, the total increase in the OASDHI rate has added a quarter of a percentage point to the NAIRU. I find little evidence of the *level* of payroll taxes having any effect on the NAIRU.

OASDHI rate increases typically take effect on January 1, making the change in OASDHI highly seasonal. The supplements component of the Productivity and Cost measure of average compensation is not seasonally adjusted, so this effect also appears in the dependent variable prior to 1980. After that, seasonal effects are removed from the Employment Cost Index. To maintain comparability, I seasonally adjust the level of OASDHI from 1980:1, using the same procedure (X-11, multiplicative adjustment) used to adjust the Employment Cost Index. This simple procedure maintains consistency between the left and right hand sides of the equation and gives plausible results. In contrast, other researchers either fail to make such an adjustment or follow procedures (such as averaging or adding seasonal dummies) the motivation of which is unclear.

2.3 Excluded effects

Table 1-3 noted a number of variables that are omitted from my wage equation. I comment on some of these excluded variables below. I discuss exclusions relating to elements of the social safety net in Chapter 3.

2.3.1 Downwards nominal wage rigidity and other non-linearities

As indicated in Section 1.2.2, my wage equation fails a Ramsey RESET test at a confidence level of 99%. This is a conventional test of functional form, including the squared fitted values in the equation. The main reason for this rejection is that the test effectively removes the constraint that the inflation coefficients sum to unity. When this constraint is not imposed, the p-value on the RESET test rises to 24%, implying that the main relationships described by the wage equation are close to linear. Nevertheless, there is evidence of more specific non-linearities not detected by this general test.

Akerlof, Dickens and Perry (1996), among others, have argued that nominal wages are rigid downwards and hence that reducing inflation will distort relative prices. This would imply wage growth would be less responsiveness at low levels and that the distribution of individual wage changes would be truncated from below at zero. As the degree of truncation increases with lower wage growth, the mean is pushed up.

To see whether some extra force is holding up wages at low levels of wage growth, I generate a dummy variable equaling 1 in the 23 (out of 200) quarters when fitted wage growth falls below 0.75 per cent (3 per cent annualized, corresponding to price inflation of under 2 per cent). When reinserted in the equation, this dummy is highly significant, with a p-value of 0.1%, boosting actual wage growth by 0.14 per cent a quarter.

Much of the explanatory power of this dummy comes from the last 12 quarters, 11 of which involve low wage growth. However, this episode by itself does not account for the result. If the equation is estimated to 1995:2, the coefficient on the low wage dummy is much the same, boosting quarterly wage growth by 0.12 per cent – though without the recent information this estimate is not precise: the p-value is 8%.

It seems then that the behavior of aggregate wages is consistent with the hypothesis of Akerlof, Dickens and Perry (1996). However, there are many possible interpretations of this result. For example, it could be due to misspecification of the functional form of another regressor. Given this, and the controversial nature of downwards nominal wage rigidity, it seems conservative to exclude this non-linearity from my basic specification. Exclusion reduces the coefficient on the minimum wage.

Another non-linearity with policy implications is the functional form of the unemployment rate. From the early studies of Phillips (1958) and Samuelson and Solow (1960), the relationship between unemployment and wage growth or inflation was viewed as convex to the origin. That is, inflation was more responsive to booms than to recessions. So, for a given average unemployment rate, there would be less inflation if unemployment was stable. Recently however, Eisner (1997) and Stiglitz (1997) have suggested that the relationship between unemployment and inflation bends in the opposite direction.

In my data set, the relationship between unemployment and wage growth is more consistent with the views of Eisner and Stiglitz, than Phillips and Samuelson-Solow. Splitting the demographically adjusted unemployment rate at its average level results in coefficients (x100) of -0.109 when above average and -0.086 when below. That is, wage growth is more responsive to recessions than to booms. This 25% difference is significant in statistical terms with a p-value of 1%, and could also be economically important. Consistent with this, entering the unemployment rate as its reciprocal or logarithm *worsens* the fit of the equation.

However, evidence of a convex Phillips curve seems fragile. The square of the unemployment rate is negatively signed, as Eisner and Stiglitz suggest, but insignificant, with a p-value of 89%. Splitting the unemployment rate at the NAIRU is also insignificant, with virtually identical coefficients.

2.3.2 Duration of unemployment

Layard, Nickell and Jackman (1991) argue that long term unemployment increases the NAIRU and hence that labor market programs should aim at reducing the average duration of unemployment. They suggest that the long-term unemployed are less effective at filling vacancies than the short-term unemployed and hence exert less downwards pressure on wages.

I can find little evidence that this effect is relevant to wage determination in the United States. As shown in Table 1-3, neither the mean duration of unemployment nor the proportion of long term unemployment have explanatory power in the wage equation. The table below estimates the different effects of short and long term unemployment when the unemployment rate is split in two. The demographic adjustment to unemployment is constrained to zero but there are no other changes to the specification. Standard errors are in square brackets.

Table 2-1: Effect of unemployment duration on wages

	Coefficient on short term unemployment rate x100	Coefficient on long term unemployment rate x100	p-value of difference
Split at 5 weeks	.184	.091	25%
	[.058]	*[.035]*	
Split at 15 weeks	.142	.091	50%
	[.036]	*[.052]*	
Split at 26 weeks	.124	.117	94%
	[.029]	*[.079]*	

Distinguishing between those unemployed for greater and less than 26 weeks seems relevant to the effects emphasized by Layard, Nickell and Jackman such as deterioration of skills, discouragement and so on. However, as the bottom cells indicate, the difference in coefficients with this split is negligible. A more dramatic contrast occurs if unemployment is split at 5 weeks. As indicated in the first row, the very recently unemployed appear to have twice the effect on wages as those unemployed for more than 5 weeks. A possible interpretation of this is that the flow into unemployment matters as well as the stock – though this might have already been inferred from the change in unemployment, which remains highly significant. These differences are in the direction suggested by Layard, Nickell and Jackman, however, none is statistically significant. The conclusion suggested by these results is that the long-term unemployed exert about the same downward pressure on wages as the short-term unemployed.

2.3.3 Variable degree of inertia

When inflation increases, so do the costs of not adjusting prices and wages. Thus optimizing models of price and wage adjustment often predict that the frequency of adjustment will increase with the inflation rate. Arguments along these lines, and supporting evidence, are presented in Ball, Mankiw and Romer (1988), Cecchetti (1984), Card and Hyslop (1997) and Taylor (1998).

The frequency of wage adjustment can be modeled in different ways. A simple approach allows the effect of prices to flow into wages more quickly in times of inflation, while leaving the total effect independent of the rate of inflation. Institutionally, such an effect is consistent with the spread of cost-of-living-adjustment (COLA) clauses in union contracts in the high-inflation 1970s and their decline in the low-inflation 1980s and 90s.

This relates to the following part of my model:

$$\Delta\text{Wage} = \psi\,\pi_1 + (1-\psi)\,\pi_{2\text{-}5} + \text{other} \qquad (6)$$

where π_1 represents last year's inflation (at a quarterly rate) and $\pi_{2\text{-}5}$ represents the average of the previous four years' inflation. The coefficient ψ reflects how quickly prices flow into wages. It can be interpreted as the speed of indexation or as the degree of inertia in the inflationary system (though further inertia arises from lags in the flow of wages into prices). In my estimated model ψ is a constant, 0.53, with standard error 0.05. However, the speed of indexation could alternatively be modeled as a function of the inflation rate:

$$\psi = \beta + \gamma\,\pi_1 \qquad (7)$$

Substituting (7) into (6) maintains the *size* of the effect of inflation (constrained to equal one), but makes the *speed of adjustment* a function of the inflation rate.

However, results of estimating this specification imply this effect is unimportant:

$$\Delta\text{Wage} = \underset{[.084]}{[\;.51} + \underset{[4.9]}{1.8}\,\pi_1]\,\pi_1 + (1 - [\;\underset{[.084]}{.51} + \underset{[4.9]}{1.8}\,\pi_1])\,\pi_{2\text{-}5} + \text{other} \qquad (8)$$

Standard errors are in brackets under the relevant coefficient. The coefficient γ is insignificant in both statistical and economic terms. The null hypothesis that the degree of inertia is independent of the rate of inflation is easily retained, with a p-value of 71%. The coefficient ψ varies in a narrow band, ranging between 0.50 (when inflation was slightly negative in 1949) to 0.56 (at 2.6 per cent quarterly inflation in 1974). Variations in the speed of indexation do not appear to be important in the US labor market.

2.3.4 Lagged Levels of Wages

Many wage equations, including those in the FRB/US model (Brayton and Tinsley, 1996), Gordon and Franz (1993), Fair (1994), Blanchard and Katz (1997), Holden and Nymoen (1998) and Sargan (1964) include the lagged level of real wages. This is often divided by productivity and interpreted as the lagged wage share or the markup of prices on unit labor costs.

If the coefficient on real wages in the wage equation is negative, then, in steady state, the real wage (or wage share) will be negatively associated with the unemployment rate. Higher real wages put downward pressure on nominal wage growth, permitting lower unemployment. This relationship can be combined with the steady state relationship implied by the price equation, (equation 2 of Chapter 1), in which the real wage and the unemployment rate are positively associated. Figure 2.3 below shows these two steady state conditions, labeled wage stability and price stability respectively. The steady state real wage and unemployment rate are simultaneously determined in both labor and product markets.

Figure 2.3: Real wages and unemployment

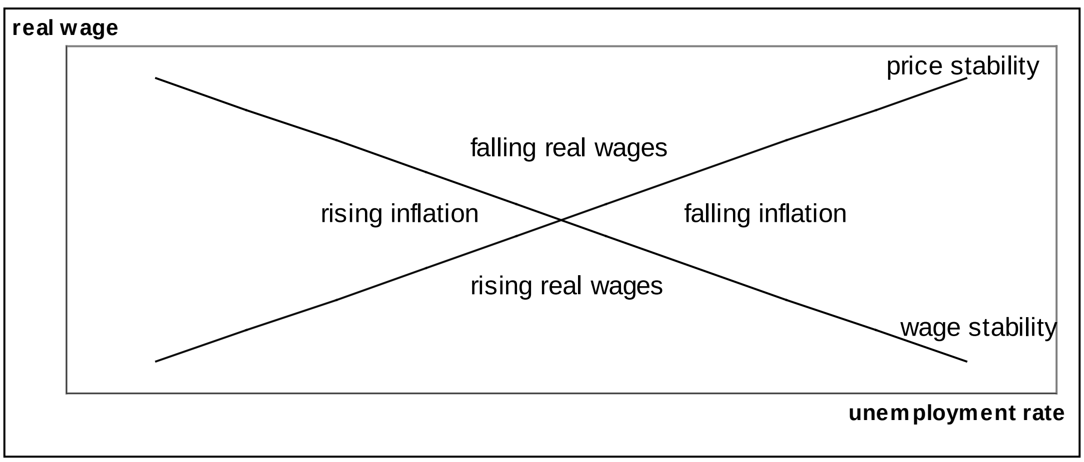

The steady state is given by the intersection of these loci. Deviations either up or down are inherently self-correcting. Deviations to the left or right might be offset by inflation-stabilizing monetary policy. In this model, product market shocks, such as increases in the price of tradable goods or the cost of capital (in the notation of Section 1.2.1, an increase in *Z*), would affect the equilibrium level of unemployment by shifting the price equation locus to the right. (Though, in the USA, shocks like these tend to be transient).

The model I outline in Section 1.2 differs from that above by excluding the lagged levels of real wages and productivity from the wage equation, making the wage stability locus in Figure 2.3 vertical. This simplifies the exposition. It implies that the wage equation determines the equilibrium unemployment rate (the NAIRU) and, given this, the price equation determines factor shares. The effect of product market shocks on the NAIRU is confined to their effect through the price wedge.

To assess this simplification, Table 2-2 shows the results of eight specifications, each of which include lags of real wages and productivity in my wage equation. Real wages are measured alternatively using the deflator for consumption expenditure and for business output (excluding agriculture and housing). Productivity is measured alternatively as its actual level or as the prediction of a linear trend with a kink at 1973. Each specification is estimated entering both variables separately and constraining the coefficients to have equal but opposite signs.

Table 2-2: Exclusion restrictions: lagged real wages and productivity

Excluded Variables	**P-value of exclusion**	**Effect on Minimum Wage**	
		Coefficient x100	**standard error x100**
I. unconstrained			
real product wage and trend productivity	6.5%	0.78	0.17
real product wage and actual productivity	3.7%	0.76	0.17
real consumption wage and trend productivity	6.1%	0.85	0.24
real consumption wage and actual productivity	3.6%	0.90	0.26

II. coefficients constrained to sum to zero

real product wage and trend productivity	21%	0.81	0.16
real product wage and actual productivity	12%	0.82	0.16
real consumption wage and trend productivity	3.7%	1.03	0.19
real consumption wage and actual productivity	1.6%	1.03	0.17

An important effect of lagged levels of real wages is difficult to discern in the US wage data. This is in contrast to European wage equations (see, for example, Holden and Nymoen, 1998). The formulations above with the greatest explanatory power trend over time. In principle, this means higher Dickey-Fuller critical values would be more appropriate than those assuming normality. In practice, such a relationship is likely to be fragile, disappearing when other trending regressors are included. Stationary formulations have little explanatory power. For example, a one standard deviation increase in the wage share (the third last row) implies a 0.04 percentage point reduction in quarterly wage growth. This reflects both lack of variation in the wage share and a small coefficient (-0.03). If the wage setting locus in Figure 2.3 does slope downward, it is quite steep.

Considerations of parsimony, ease of exposition and conservatism suggest excluding the levels of real wages and productivity. Nevertheless, a more elaborate model that sought to estimate product market effects, or encompass the literature on European unemployment, could include such a term. Furthermore, inclusion facilitates one particular theoretical interpretation I develop in Chapter 4.

2.4 Confidence intervals

It may be reasonable to assume that the coefficients in my wage equation are normally distributed, though I discuss some problems with this assumption below. However, estimating effects on the NAIRU involves taking ratios of these coefficients. For example, the effect of the minimum wage on the NAIRU is given by dividing the coefficient on the minimum wage in the wage equation by that on unemployment. These ratios will not be normal. The ratio of two independent standard normal variables has a Cauchy distribution, which has fatter tails than the normal. If the denominator has a non-zero mean, the Cauchy distribution will be non-centered and skewed. Correlation between the numerator and denominator imparts further skewness. Intuitively, near-zero realizations of the denominator result in unusually dispersed realizations of the ratio.

To allow for these effects, I construct confidence intervals about two of my estimates: the most recent (1998:2) level of the NAIRU and the effect of the minimum wage on the NAIRU. The procedure is as follows. I start with the coefficients and their covariance matrix estimated in the wage equation. I then take 10,000 random draws of a vector of independent standardized normal variables. Multiplying each of these by the Cholesky factorization of the covariance matrix and adding the coefficient vector creates 10,000 draws of a vector that has a multivariate normal distribution with the same mean and covariance as the coefficients in the wage equation. From this, I calculate 10,000 realizations of the coefficients in equation (8) of Chapter 1, including the effect of the minimum wage on the NAIRU. Taking the 1998:2 values of the variables in equation (8) as given, I also calculate 10,000 realizations of the NAIRU for this period. The results are shown in Figure 2.4 below. A 95% confidence interval for the NAIRU in 1998:2 is (4.8, 5.6); for the minimum wage coefficient it is (3.6, 9.0).

Figure 2.4: Confidence intervals

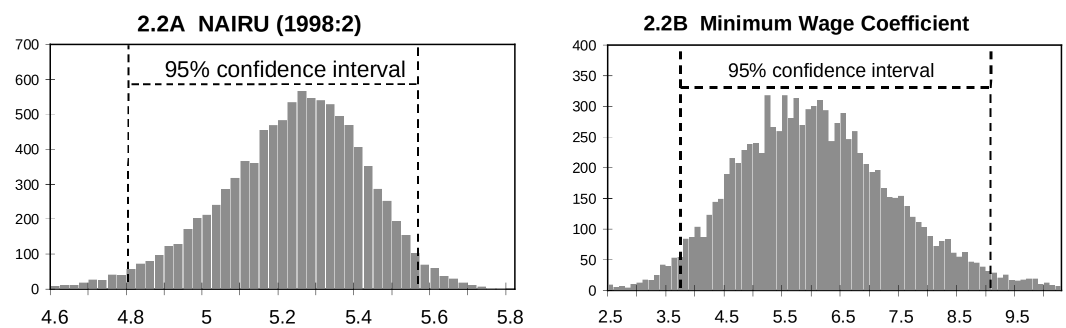

This approach assumes that the least squares coefficients in the wage equation have a multivariate normal distribution. So does my presentation of many hypotheses as p-values. But if regressors are dynamic, as they are in my wage equation, then normality may not apply. The central limit theorem applies asymptotically when regressors are stationary.

Consider the relative minimum wage (though similar considerations apply to the replacement rate). This is a ratio of two integrated variables: the coverage-adjusted minimum wage and the average wage level, both in nominal terms. The simplest case to consider is when the minimum wage adjusts so as to maintain proportionality with the average wage. Then the two variables would be cointegrated. The relative minimum wage would be stationary and conventional asymptotic approximations would apply.

Alternatively, suppose that the nominal minimum wage is strongly exogenous. A similar argument applies if the coefficient on the relative level of the minimum wage in the wage equation is positive. This implies that the average wage "error-corrects" to the minimum wage and that the two variables are cointegrated. Under this "alternative hypothesis", the regressor is stationary and normality is again a reasonable approximation.

However, the most interesting hypothesis to be tested is that this coefficient is zero rather than positive. If the nominal minimum is assumed to be strongly exogenous then this null hypothesis implies that the minimum wage and the average wage are not cointegrated. Kremers, Ericsson and Dolado (1992) show that "error-correction" coefficients in conditions like this have an asymptotic distribution under the null that lies between a normal and Dickey-Fuller distribution. However, the distribution will be approximately normal when, as is the case here, the impact effect of a change in the minimum wage is much less than its long run effect and the variance of the change in the minimum wage greatly exceeds that of the residuals in the equation. Kremers, Ericsson and Dolado's Monte Carlo experiments suggest that conventional Gaussian critical values will tend to over-reject the null hypothesis. However, extrapolating their results to my parameter values ($a = 0.03$, $s = 17.4$, $q = 16.9$) suggests this bias is small. Normality seems a reasonable approximation, even under the null hypothesis.

In short, whatever assumptions are made about the underlying regressors, it seems reasonable to assume that the coefficients in my wage equation are normally distributed. From these it is possible to construct confidence intervals about estimates of the NAIRU, which are slightly non-normal.

2.5 Heteroskedasticity

Data quality improves over time as statistical bureaus refine their techniques, attract greater resources and exploit economies of scale (such as the law of large numbers). Reductions in measurement error affect most macroeconomic series, however they are particularly important in the wage data. This is evident in Figure 2.5, which shows the residuals and standard error of my wage equation. The standard error is estimated to shrink from 0.56 percentage points in 1948 to 0.12 percentage points in 1998 – a rate of decline of 3.1 per cent a year.

Figure 2.5: Residuals and standard error of equation

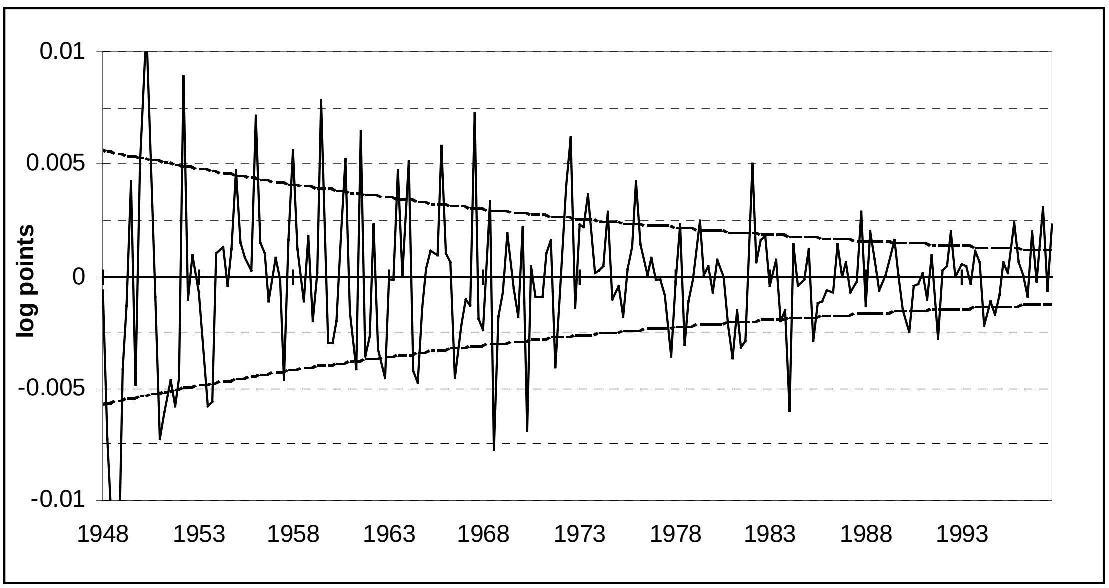

When estimated by unweighted least squares, my baseline equation exhibits strong heteroskedasticity. A Breusch-Pagan test (Greene p552) regressing the squared residuals on a constant and linear trend rejects the hypothesis of homoskedasticity with a p-value of 0.0002%. The frequent large residuals at the beginning of the sample are also reflected in fat tails of the distribution. A Jarque-Bera test rejects the hypothesis of normality with a p-value of 0.001%.

If an equation is estimated by least squares under the invalid assumption of homoskedasticity the estimated coefficients will be inefficient (albeit consistent – if the other classical assumptions hold) and estimated standard errors will be inconsistent. Consistent standard errors can be estimated by the "jackknife" procedure of MacKinnon and White (1985) or the simpler measure suggested Davidson and MacKinnon (1993, Ch16.3). However, the coefficients remain inefficient. Essentially, later data are more informative and so should receive a higher weight. Furthermore, the non-normality of the residuals makes inference difficult.

As discussed in Section 1.4.1, my estimates of the effect of the minimum wage do not change much over time. Accordingly, how, (or whether) a time-varying variance is dealt with will not affect this qualitative result. However, for many purposes quantitative precision is required and so efficient estimation techniques should be used.

A simple way of allowing for improved data quality is two-stage weighted least squares, or feasible GLS. The squared residuals (or their logarithms, to keep the predicted variance positive) from a least squares equation can be regressed on time, with the fitted values from this auxiliary regression being used as weights in a subsequent regression. While an improvement on OLS, feasible GLS remains inefficient. If the residuals are normally distributed, the squared standardized residuals will be Chi-squared. OLS on a Chi-squared distributed variable, or its logarithm, is asymptotically inefficient, in contrast to maximum likelihood, which converges to the

Cramer-Rao minimum variance. To be precise, Harvey (1976) shows that the variance of estimates from this two-step procedure is 2.5 times as large as that of maximum-likelihood estimates. Furthermore, the estimate of the equation standard error is inconsistent and biased downward. Because the bias is a constant proportion, this does not affect the relative weighting of least squares estimates. However, it does make drawing inferences about the fit of the equation difficult.

Consistent, asymptotically efficient, maximum likelihood estimates are relatively easy to compute. As Harvey (1976) and Greene (1996 p567 – though beware of typos) describe, numerical maximization of the likelihood involves iterating on two equations.

At the *i*'th iteration the weighted least squares parameters equal

$$\beta_{i+1} = \left(X'\Omega_i^{-1}X\right)^{-1} X'\Omega_i^{-1} y \qquad (9)$$

where the diagonal elements of Ω_i are the predicted variance (σ_{ti}^2, described below) of each observation. This gives residuals $\varepsilon_{ti} = y_t - X_t'\beta_i$. Given these residuals and the predicted variance, a vector U_i can be formed, with elements comprising $\frac{\varepsilon_{ti}^2}{\sigma_{ti}^2} - 1$. If the variance declines exponentially over time, then $\sigma_{ti}^2 = e^{\gamma_i + \delta_i TIME} = e^{\alpha_i Z}$ where Z is the matrix [1 TIME]. The coefficients in α are updated by

$$\alpha_{i+1} = \alpha_i + (Z'Z)^{-1} Z'U \qquad (10)$$

The asymptotic variance matrix of $\alpha = 2(Z'Z)^{-1}$.

For initial values, $\alpha_1 = 0$ can be used, in which case $\beta_1 = \beta_{OLS}$. If, alternatively a consistent estimator of α were used (Harvey shows how this might be computed), then only one iteration would be needed to produce an estimator with the same asymptotic distribution as the maximum likelihood estimator.

Heteroskedasticity is to be expected – the rationale for splicing the compensation measure from the Productivity and Cost release with that from the Employment Cost Index is to reduce measurement error. But might it not also reflect some economic process that should be modeled directly? It is easy to speculate how trends such as declining unionization or improved access to macroeconomic forecasts might explain the heteroskedastic residuals. However, to the extent that such fundamental forces are captured by an exponential trend, they would not affect the estimates.

In any case, this interpretation seems unlikely. We can conclude that innovations in measurement, such as the introduction of the Employment Cost Index in 1980, have reduced the noise in the data, by comparing it with the previous series, which has been maintained. When my equation is estimated with unweighted least squares then using the initial measure of compensation for the whole sample period raises the equation standard error by 24 per cent (admittedly, on a specification designed to describe the spliced series, except that I no longer seasonally adjust payroll taxes) and tests for heteroskedasticity are no longer significant. More tellingly, if compensation per hour is substituted for the spliced wage series in my price regression (Appendix 1 of Chapter 1) the standard error of the equation increases by 17 per cent, and the coefficients on the three unit labor cost terms each decrease while their standard errors increase. This implies that splicing the wage data results in a series that is both more accurate and more useful.

Surprisingly however, the variance in the data does not exhibit discrete reductions at the time of such innovations, but appears to gradually decline over time. My attempts to model the heteroskedasticity with dummy variables reflecting known data-breaks were unsuccessful. A break in the intercept of the variance equation at 1980:1 has a coefficient (0.33, with standard error 0.37) that is small, *positive* and statistically insignificant. The break *increases* the coefficient on time. Gradual improvements arising from learning by doing and economies of scale appear to be important. I did not explore stochastic trends, beyond noting that estimating the model as a GARCH(1,1) does not affect the coefficient on the level of the minimum wage or its standard error.

Weighting the data affects the validity and robustness of the estimates rather than their magnitudes. (Not surprisingly, given that most coefficients are stable). The main changes are a 50 per cent increase in the coefficient on the unemployment rate (partly offset by a reduction in the coefficient on the change in unemployment) and a large change (from near zero) in the coefficient on unemployment benefits.

2.6 Overall stability of the wage equation

The overall stability of a heteroskedastic equation at a given breakpoint can be assessed using a Wald test (Greene p354). This a generalization of the familiar Chow F-test, simplifying to the latter under homoskedasticity. The gray solid line in Figure 2.6 below shows the sequence of Wald tests for breakpoints from 1953 to 1995. The test statistics can be interpreted as the probability of obtaining different coefficients before and after any specific breakpoint if the null-hypothesis of no structural change is valid.

Figure 2.6: Probability of no structural break

Wald breakpoint tests on equations with and without minimum wage

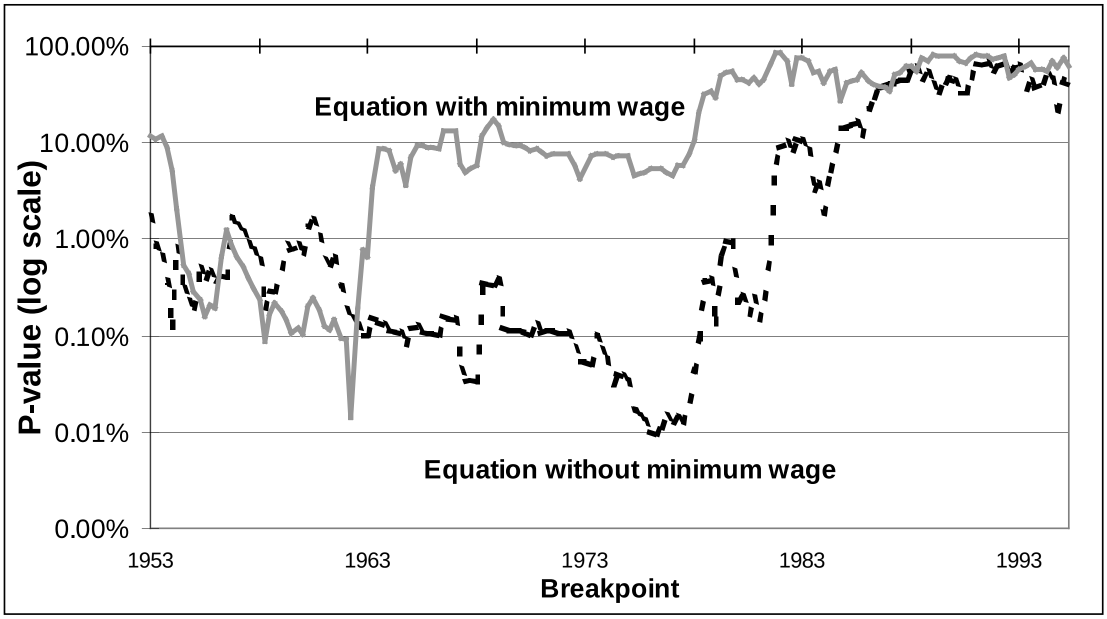

If the breakpoint is not known, the overall significance of the sequence of tests can be gauged by a procedure discussed in Andrews, Lee and Ploberger (1996). This involves calculating the weighted average of the Wald Chi-squared statistics across a sequence of breakpoints. This statistic is then compared with critical values tabulated by Andrews and Ploberger (1994). For breakpoints from 1963:3 (the fourth adjustment in the minimum wage) to 1995:3 (after which the regressors are collinear), this Andrews-Ploberger statistic is 7.6, below the 10 per cent critical value of 8.1 (for Andrews-Ploberger's parameters of p= 9 and λ = 7). Thus, apart from the early subsamples discussed in Section 1.4.1, the overall stability of the equation seems satisfactory.

In contrast, an equation estimated without the level of the minimum wage shows much stronger evidence of instability, despite its weaker statistical power. The dashed black line in Figure 2.6 shows Wald breakpoint tests for the equation when it is estimated omitting the level of the minimum wage. An Andrews-Ploberger statistic for 1963:3 to 1995:3 is 13.6, exceeding the 1 per cent critical value (p=8, $\lambda = 7$) of 10.6. (Overall stability is still rejected at the 1 per cent level if all breakpoints prior to 1978 are disregarded).

One aspect of this instability has recently attracted attention. An equation that does not control for the minimum wage shows a large unexplained reduction in the NAIRU over the last two decades. A split in the intercept at 1982:4, (when the coverage adjusted relative minimum wage reached its lowest point for 20 years) implies a reduction in the NAIRU of 1.4 percentage points with a p-value of 0.8%. However, when the level of the minimum wage is included, this otherwise unexplained break disappears – the intercept term in the NAIRU *increases* by 0.2 percentage points, with a p-value of 63%. As Hendry (1980 p275) argues, "an essential prerequisite for tentatively accepting a model is that it can explain previous findings, including why such models broke down when they did." A wage equation that includes the minimum wage "encompasses" previous research in that it explains why the NAIRU estimated by simpler wage equations is unstable.

2.7 Existence of the NAIRU

The existence of a unique NAIRU is difficult to establish with my data set. When unrestricted, the coefficients on the lagged inflation terms sum to 0.82, with a standard error of 0.06. The hypothesis that they sum to 1 is rejected with a p-value of 0.2%. Experimenting with alternative lag lengths and specifications produced similar results. For example, a general, over-parameterized model with 30 lags lying on a 4^{th} order polynomial also sums to 0.82, with a standard error of 0.06. The sum of coefficients declines slightly for lags of less than 4 years.

If these coefficients were assumed to be stable they would imply the existence of a long-run steep "tradeoff" between inflation and unemployment. That is, a reduction in unemployment will give rise to higher, but not constantly rising inflation. There is no unemployment rate below which prices continually accelerate. Any unemployment rate would eventually be consistent with stable inflation.

Without constraining the inflation coefficients to sum to one, my equation still generates a highly significant coefficient on the level of the minimum wage. The effect on inflation remains important (indeed, slightly larger in magnitude) and persistent – though no longer permanent. My equation would then imply that, instead of shifting the NAIRU, the minimum wage shifts a steep long-run Phillips curve. The policy implications would presumably be similar, though more complicated.

However, my impression is that few researchers or policy makers would be interested in such an interpretation. The possibility of a long-run tradeoff is widely considered to be implausible and unrealistic. A well-defined NAIRU (that is, one that is independent of the rate of inflation) can coexist with inflation coefficients summing to less than one because inflation is expected to revert to its mean or because of errors in the measurement of prices. Constraining my inflation coefficients to sum to one gives my results a simple and important interpretation.

It does however make it harder to assess their empirical validity. The unrealistic restriction means that specification checks, such as inclusion of the squared fitted values or of the lagged dependent variable, indicate the existence of problems. The estimated effect of minimum wages to nominal wage growth seems to be robust to variations addressing these. Nevertheless, it is uncomfortable relying on a model that does not realistically describe the data.