CHANGING PUBLIC SECTOR WAGE DIFFERENTIALS IN THE UK

Richard Disney
Amanda Gosling

THE INSTITUTE FOR FISCAL STUDIES
WP08/02

Changing public sector wage differentials in the UK

by

Richard Disney[†]   and   Amanda Gosling[‡]

[†] University of Nottingham, and Institute for Fiscal Studies
[‡] University of Kent at Canterbury, and Institute for Fiscal Studies

Abstract

The paper estimates public sector wage differentials and their changes over time for men and women in the United Kingdom using the New Earnings Survey. It presents estimates that are robust to unobserved workforce characteristics and that also show the impact of policy changes and cyclical factors. The methodology also allows us to examine the extent to which discrepancies in public and private sector pay induce changing relative qualities of the sectoral workforces.

JEL Classification   J31   J58

Key words   Public sector   wages   unobservables

Changing public sector wage differentials in the UK

1. Introduction

This paper estimates public sector wage differentials in the United Kingdom (UK) using New Earnings Survey (NES)/ Annual Survey of Hours and Earnings (ASHE) panel data over the 1970s, 1980s, 1990s and early 2000s.[1] In contrast to existing studies that examine public sector pay using standard panel data methods , a central and desirable feature of our method is that it allows the public sector 'penalty' or 'premium' to be change over time while still exploiting the panel dimension of our data to control for possible unobserved differences in worker characteristics between the two sectors. It is this time variation in the "like for like" difference in pay between the two sectors that is of interest to stakeholders in the public sector wage-setting process: government departments, employers, unions and those making public sector pay recommendations such as the UK's Pay Review Bodies.

We believe that our approach to measuring a time-varying public sector wage differential also has another desirable feature: it is consistent with the existence of substantial unobserved differences in characteristics between private and public sector workers. Since we control for differences in the quality of workers, we can thereby show the impact of policy changes and cyclical factors both on relative pay and (indirectly) on relative *quality*. Thus our methodology also allows us to examine the extent to which discrepancies in public and private sector pay induce changes in the relative qualities (observed and unobserved skills) of the respective sectoral workforces.

Our data are the NES/ASHE panel data on individual hourly wages,[2] collected from employers at the workplace using a 1% sample of employees. Figure 1 plots 'raw' public-private hourly pay differentials (percentage differences at the mean) for men and women from the NES/ASHE for the period 1975 to 2006. Both the male and female pay differentials have fluctuated considerably over time, falling during the late 1970s and 1990s and rising in the mid-to-late 1980s. For example the measured female differential more than halved between the mid-1970s and the mid-1980s as did the male differential in the mid-1990s. In addition, Figure 1 shows that average public sector pay, especially for women, exceeds that of private sector workers. The aim of this

---

[1] ASHE replaced NES in the early 2000s.
[2] Measured as the reported weekly wage divided by reported weekly hours.

paper is to assess the extent to which these levels and changes in differentials are 'real' and not just a statistical artefact arising from the changing composition of the workforce. Such compositional changes may have arisen from the changing composition of public sector jobs – for example the privatisations and contracting-out of public sector jobs in the period after the mid-1980s, and also as a result of changes in the characteristics of workers that are prepared to accept jobs in both sectors.[3] Since the quality of public sector workers may be endogenous to the perceived wage differential (as in Nickell and Quintini, 2002), adjusting for labour quality is a central problem for studies of public sector pay that have used standard wage equations to estimate the average 'premium' or 'penalty' attached to working in the public sector, for which we shall use the notation 'β'.

With a time-varying wage differential between sectors, the relative gains to worker in each sector, which underlies the selection of workers into sectors, will also be time-varying. Many studies that have addressed the selection issue in the field of public sector pay have either used selectivity correction or instrumental variable methods. This necessitates the presence of something in the data set that 'explains' public sector status but which is uncorrelated with unobserved factors that explain wages. Plausible instruments of this nature are hard to find and the findings from such studies are generally inconclusive.[4] Some researchers (Disney and Gosling, 2003; Haskel and Szymanski, 1993; Monteiro, 2004) have implicitly or explicitly exploited privatisations as 'natural experiments', but estimates of β derived from such methods are typically small and poorly-defined. The problem in all such studies is that they rely on the assumption that the 'public sector pay effect' is stable over time (and, especially, over the period of the 'treatment') and also that the effect can be generalised from the 'treated' group.

---

[3] Disney, Goodman, Gosling and Trinder (1999) control for the changing occupational composition of the public sector workforce. This strategy eliminates the measured average positive pay differential for public sector men, although a (time-varying) female differential remains.

[4] Studies for a variety of countries which control for selection include those that identify off functional form, especially on the education variables, such as Belman and Heywood (1989) for the US, Stelcner, M., van der Gaag, J. and Vijverberg, W. (1989) for Peru, van der Gaag and Vijverberg (1989) for Cote d'Ivoire, and van Ophem (1993) for the Netherlands, or off background family characteristics, as in Terrell (1993) for Haiti, Borland, Hirschberg and Lye (1996) for Australia, and Dustman and van Soest (1998) for Germany. From a reading of their paper, the identification strategy of Rees and Shah (1995) for the UK is not clear.

Another common strategy (where appropriate data are available, as here) is to exploit variation over time in individual wage and public sector histories. Existing work using panel data typically utilises a 'fixed effect' model which rules out any time variation in the average estimated β, and thus any possibility that the 'true' βs can fluctuate over the sample period. This restriction is not only implausible given the time variation of the raw data in Figure 1 but also averages out the precise variation in the data that is of interest, namely the changes over time in β. It is, in fact, an unnecessarily restrictive model to use, since other, less restrictive, models have been suggested for panel data in the (closely related) literature on estimating empirically the 'mark-up' attached to being a member of a union, or covered by a union agreement.

The present paper utilises a two-step procedure for estimation of a potentially time-varying β using panel data. First we implement a new test for the null hypothesis that changes over time in wages are the same for both public and private sector workers. The yearly contributions to the test statistic provide a first stage indication of any differential cycles and trends in wages across the two sectors. In the second stage these are adjusted to take into account any possible long run divergence in pay.

In the first stage a simple "within groups" model is estimated. On the LHS is the log(hourly) wage and the right hand side includes an individual fixed effect, age and its square and a full set of time dummies. Note that public sector status is not included as an explanatory variable. The working assumption is then that any observed difference between public and private sector workers' wages is driven by composition effects and/or sampling variation. If this assumption is correct, the residuals from this equation should be uncorrelated with public sector status at any point in time, thus we have a simple test of the null hypothesis that there are no public sector effects on wages We show that this null hypothesis can easily be rejected. We conduct Monte Carlo simulations to confirm that our procedure for testing the null hypothesis is a valid one.

We then utilise the test statistics for each year to give a preliminary estimate of the public sector pay effect which is identical across the sampled individuals but varies over time. However a potential bias in the estimate of the average effect in each year arises from its failure to consider the effect of trajectory of the individual's public sector status over the whole time period on her average wages. Nevertheless we can show that the first stage test statistic for each year together with the individual histories of wages

and public sector status give us a set of T (T=number of years of data) linear simultaneous equations that allow us to correct these potential biases.

We illustrate the results of this estimation method on the NES data set for the public sector as a whole for men and women. Given the size of the data set, we are then able to illustrate the implementation of our methodology for a particular sub-group of workers. We focus on nurses, whose pay bargaining since the mid-1980s has been covered by a Review Body procedure and for whom have been several pay restructurings over the period. We wish to see whether such pay policies have affected the differential (if any) attached to nurses' pay over time. We compare the pay of public sector nurses with a control group of workers from the data set who are identified as having *ever* been nurses over the whole period. We can thereby show not just how pay policies impact on the time variation in nurses' pay relative to the control group, but also how changes in relativities over periods of several years have implications for the composition of the public nursing workforce. We discuss the implications of these results for pay policy. We summarise the paper in the concluding section.

2. Estimating pay differentials from panel data: all public sector workers

2.1. Evidence from standard panel data models

Consider the following model of sectoral pay determination:

$$w_{it} = \beta_t P_{it} + \alpha_t x_{it} + \eta_{it} \tag{1}$$

In equation (1), w_{it} denotes the (log) hourly wage of individual i at time t, P is a dummy variable denoting sectoral status ($P=1$ for public sector status), x is a vector of control variables such as age and qualifications and η is an error term which includes factors such as unobserved skill/ability, α and β are coefficient estimates. Least squares estimation of (1) is likely to yield inconsistent estimates as it is likely that η and P will be correlated given the selection of workers into sectors (Disney and Gosling, 2003).

To deduce anything about β, therefore, we need to place some structure on (1). In an error components model, write:

$$\eta_{it} = f_i + \varphi_t + \varepsilon_{it} \tag{2}$$

where η is the sum of an individual person effect assumed constant over time, f_i, a macro shock assumed constant over individuals, φ_t and an innovation ε_{it} assumed unrelated to the individual person effect, to public sector status and to β. A standard strategy (fixed effects estimation) places the restriction on (1) and (2) jointly that β is constant over time, although the person effect can be augmented by an individual-specific growth term, as in (3):

$$w_{it} = \beta P_{it} + \alpha_t x_{it} + f_i + g_i t + \varphi_t + \varepsilon_{it} \qquad (3)$$

As a benchmark for our own, alternative, subsequent method we estimate the standard model in (1a) for male workers using the NES/ASHE data set from 1975-2006 with just over 120,000 person-observations in each year. We regress the log of real hourly earnings on the person effect, f_i, a quadratic in age, public sector status in each year, and time dummies. We derive three specifications for men: a fixed person-effect estimate, a person-effect specification augmented by an individual earnings growth factor g, and a person effect specification estimated over a rolling five-year 'window' to obtain some time variation in the average estimated βs.

The results are depicted in Figure 2. The first two specifications force β to be time invariant, with the fixed person effects suggesting a small public sector penalty whereas the specification augmenting this with an individual growth factor implies a small public sector premium. The five year window estimates do give time-varying βs, with the results of this exercise tracking the decline in the 'premium' in the first part of the period and suggest an increasing 'premium' in the early 1990s and the early 2000s. Note however that the raw data in Figure 1 and the moving person effect estimates in Figure 2 give a total different trajectory for the period post-1995 and that the latter estimate gives very little volatility in the mid-1980s.

2.2. An alternative approach for time-varying βs

We adopt an alternative approach to this standard framework. Ignoring the individual growth effects, we start by testing the null hypothesis that β is zero for all t. Then equation (3) collapses to:

$$w_{it} = \alpha_t x_{it} + f_i + \varphi_t + \varepsilon_{it} \qquad (3a)$$

which can again easily be estimated by fixed effects or a within-groups estimator. This will give us consistent estimates of ε_{it} *under the null*. This also implies that:

$$E(\hat{\varepsilon}|time, x, P = 1) = E(\hat{\varepsilon}|time, x, P = 0) = 0 \qquad (4)$$

We again implement the model using the NES/ASHE data set from 1975 to 2006. We estimate equation 3(a) by regressing the log of real hourly earnings on the person effect, f_i, a quadratic in age and time dummies. We then use the residuals from the estimated equation to provide a straightforward test that $\beta = 0$ for all t. If the squared differences in the estimated innovations between public and private sector workers exceed some critical value, we can thereby reject the null in (4). A natural test statistic of the null is the sum of the squared yearly residuals. We obtain the distribution of this test statistic and its critical value under the null by re-centering the bootstrapped standard errors at zero.[5] The comparison of the calculated values of the test statistic and the critical value of the test are provided in Table 1, and this table shows that the null hypothesis can easily be rejected – in other words, in at least some periods, β, is significantly positive or negative.

Table 1: Testing the null hypothesis that β = 0

	Test statistic	Critical value at 5% level	N
Men	780.75	22.02	1859964
Women	839.96	22.67	1401643
All workers	1620.71	41.63	3261607

Note: N= number of person years (no. of years = 27)

The *yearly* values of the test statistic provide a good indication of how the public sector wage differential for particular groups may be changing over time. In effect, the test statistic for each year measures the component of the residual variance that can be 'explained' by the sectoral affiliation of the worker in that year. The value of this test

---

[5] The precise details of this procedure can be obtained from the authors on request. We have not yet worked out the complete analytical properties of this test, but we have undertaken Monte Carlo simulations to investigate the rejection probabilities when the null hypothesis is true. We find that our test rejects less frequently: for example at the 5% level our test rejects only 3% of the time. The tests in Table 1 are in fact for the NES data 1975-2001; we are updating to 2006.

statistic can therefore provide a first approximation of the public sector 'effect' on wage at a point in time, under the assumption that this predicted effect, which we denote as $\hat{\beta}$, is time-varying but constant across individuals in the particular sample (which can of course be a sub-sample of the total population differentiated by some characteristic such as sex). We therefore term this, in subsequent Figures, the 'first stage' estimate of the public sector pay effect.

The method used to calculate what we then term the 'second stage' adjusted prediction is related to an idea suggested by Jakubson (1991) in the context of his discussion of alternative panel methods that might be developed to measure the trade union pay effect over and above a simple 'fixed effects' estimator. As in Jakubson's paper, our method involves consideration of the sectoral status of the individual not just in the current period but in *all* periods of the data set, and we show that the failure to consider the duration and timing of public sector status of individuals may bias the first stage estimator that we have proposed. We continue, however, to assume that the 'person effect' is constant over the period and that the same (time-varying) public sector pay effect applies to all the sampled individuals.

To illustrate the second stage adjustment of our 'first stage' predictions of the year-on-year public sector pay effects, we consider a very simple model in which we observe individual *i* only in two periods, 1 and 2. Netting out the person effect and the covariates, the mean (log) wage of individual *i* is determined by the probability of being in the public sector in each of periods 1 and 2 (P_1 and P_2), times the respective 'true' public sector pay effects in each period, β_1 and β_2, as illustrated in (4).

$$\overline{W}_i = \frac{1}{2}[\beta_1 P_{1i} + \beta_2 P_{2i}] \tag{5}$$

Write the deviation of the wage in period 1 from this mean wage as:

$$W_{1i} - \overline{W}_i = \beta P_{1i} - \frac{1}{2}[\beta_1 P_{1i} + \beta_2 P_{2i}] \tag{6}$$

We now demonstrate how our first stage estimate of the public sector pay effect in any period $\hat{\beta}_t$ is potentially biased and how we can correct this bias. Write the

expected measured residual R in period 1 given that the individual is observed to be in the public sector as:

$$E(R_{1i}|P_{1i}=1)=\beta_1-\left[\frac{1}{2}\beta_1 E(P_{1i}|P_{1i}=1)+\frac{1}{2}\beta_2 E(P_{2i}|P_{1i}=1)\right]$$

$$=\frac{1}{2}\left[\beta_1-\beta_2 E(P_{2i}|P_{1i}=1)\right]$$

(7a)

That is, the residual is the within-period 1 true public sector 'effect' relative to the effect of being in the public sector in all or any other periods. Similarly, define the residual within period 1 from *not* being in the public sector as:

$$E(R_{1i}|P_{1i}=0)=-\frac{1}{2}\beta_2 E(P_{2i}|P_{1i}=0)$$

(7b)

We can therefore define our 'first stage' measured public sector effect as the difference in these residuals relative to the mean [i.e. (7a) – (7b)] as:

$$\hat{\beta_1}=\frac{1}{2}\beta_1+\frac{1}{2}\beta_2[E(P_{2i}|P_{1i}=0)-E(P_{2i}|P_{1i}=1)]$$

(8)

Note that the bias in equation (8) will be smaller the lower the correlation between public sector status over time and larger when the βs are all of the same sign. By analogy, we can construct a similar equation describing the measured public sector effect in period 2, $\hat{\beta_2}$, also as a product of two unknowns, the 'true' β_1,β_2. So we have two (linear) equations in the predicted $\hat{\beta}s$ to explain the two unknown βs. In this way we can solve for the unknown *β*s. These solved-out *β*s are our 'second stage' estimates of the public sector wage effect. It is also important to emphasise that this second stage estimator permits the average measured public sector pay effect over the whole period to deviate from zero, as well as in any given year. Again we bootstrap to provide a measured 95% confidence interval round the estimates.

Figure 3 Panel A and Panel B illustrate the results for public sector men and women respectively of our first and second stage estimates of the pay differential. The 95% confidence intervals derived from bootstrapped standard errors are tightly defined

around the estimates.[6] At first sight, the point estimates look like the raw differentials in Figure 1 but there are two key differences.

The first is that the average differential between public and private sector pay is close to zero for both men and women – this is not surprising for the first stage estimates since we are utilising the residuals from the fitted model in equation (4) but it is not imposed at the second stage. Indeed we can see that the second stage estimation procedure slightly increases the public sector 'penalty' for men (as expected, since this replicates the fixed-person estimate in Figure 2) whilst introducing a small premium for public sector women. In fact, using the confidence intervals around the second stage estimates, the second stage-estimate for men of the public sector pay effect is significantly negative in 20 of the 32 years of data, while for women it is significantly positive in 18 of the 32 years of data (but significantly negative in 8). We do not find evidence of the sizeable positive premium for public sector women found in other studies, including our own (Disney and Gosling, 1998), which use standard estimation methods. That the long run public-private pay differences is not very large is consistent with a number of 'stories': see for example, Postel-Vinay and Turon (2007).

The second difference is that, while at first sight, the variability of adjusted public pay differentials for men and women seem to reflect the raw differentials in Figure 1, there are important disparities between Figure 1 and Figure 3. For example, we do not observe in Figure 3 the trend decline in public sector pay relative to private sector pay between 1975 and 1987 that is observed in the raw data, but rather a broadly stable small difference other than a significant adverse 'shock' in 1979. Both the raw data and our method suggest a boom in public sector pay, relative to private, in the mid-1990s, although our data suggest that the peak is slightly earlier than in the raw data. Both the raw data and our results agree that public sector pay declined from the mid-1990s onwards; contrast this with the results in Figure 2 for the 5-year rolling average of the standard person effects model. In fact, pooling our estimates over the time period suggests weak evidence of counter-cyclicality in the public-private sector pay differential: the coefficient of a regression of the estimated second stage public pay effect on time, the unemployment rate and the inflation rate gives a just significant coefficient on unemployment rate that is positive (+0.006, s.e. 0.003). This accords with

---

[6] To avoid cluttering the Figures, we have not shown the confidence intervals; they are available from the authors on request.

a simple model of pay bargaining in which private sector wages are more sensitive to labour market conditions than public sector wages.

3. Evidence for a public sector occupation: nurses

In this section, we apply our method to examine wages in a public-sector dominated occupation: that of female nurses and midwives (hereafter referred to as 'nurses' for brevity). The dimensionality of the NES (over 3 million person-observations, see Table 1) allows us to examine some of the larger occupations and professions, and the nurses' labour market is of particular interest given the number of pay reforms and pay awards that have been targeted specifically at this group.

Buchan and Seccombe (2005) estimate that there were around 640,000 qualified nurses registered with the Royal College of Nurses in 2005, of which 460,000 were employed in the NHS, and another 16,000 in GP practices. In addition there were at least 40,000 health care assistants employed in the NHS. Some qualified nurses are either not working or working in another occupation, or working in the private sector: for example around 60,000 qualified nurses were identified in 2005 as working in private care homes. So nursing is a public-sector dominated profession, but a significant minority of nurses at any one time are not working in the public sector.

In our panel data set, we can identify whether a current worker *ever worked or ever will work* as a nurse and this permits us a comparison group with our public sector group, which is nurses employed at any point in time in the public sector. At any one time, roughly three-quarters of those who *ever* worked as a nurse are currently employed as a public sector nurse; of the remainder (the control group) a large fraction (but not a majority) are working in the private sector in residential care homes, private clinics etc as nurses. However the majority of the control group who are working are employed in a wide variety of settings including office work, management and other white-collar occupations, but also in manual jobs such as cleaning.

The nursing labour market has been subject to a series of pay regimes and pay policies over the period of our data set. As our data set starts in 1975, the first significant event is the 'Social Contract' pursued by the Labour government in the late-1970s, which appears to have had a significant effect in restraining public sector wages across the board. At the end of that Labour administration, the 'Clegg Commission' was established to consider public-private wage differentials, and it recommended

significant large 'catch-up' awards for public sector groups, including nurses, which the incoming Conservative government – somewhat reluctantly – honoured in 1980.[7] In 1983, the government decided to introduce a statutory body for making pay recommendations for nurses based on evidence provided by staff, employers and the government, with the government remaining the final arbiter of pay awards. This body, originally known as the Review Body for Nurses and Allied Professions, has gradually extended its remit; becoming the Nurses and Other Health Professions Review Body in 2004 and the NHS Pay Review Body in 2007, covering almost 1.5 million NHS workers.[8]

In 1988-89, in the face of recruitment difficulties, the government undertook a thorough regrading exercise, in which many traditional nursing grades disappeared and were replaced by a more explicit career structure. The reform incorporated substantial pay rises for some segments of the workforce, a plan for further support for nurses' training (known as *Project 2000*) and the introduction of what were subsequently to be known as High Cost Area Supplements banded in three areas around London (more commonly known as 'London Weighting'). After this time, Review Body-recommended increases kept pace with or exceeded inflation but the economic boom that commenced soon after the UK left the ERM led to a generally more rapid upsurge in private sector remuneration. Despite several major reorganisations of the NHS, the nurses' pay structure remained relatively untouched until another major review and regarding exercise was implemented in 2004-05 known as *Agenda for Change (AfC)*. AfC had several objectives, including simplification of the complicated set of pay rates and working practices that had developed since the pay reforms in the late 1980s, and in particular a desire to 'equal value proof' pay across NHS occupations with different sex compositions of the respective workforces.

To examine the effects of these changes, we examine the evolution of the pay differential between public sector nurses and our 'control' group (as defined above) over the 1975-2006 period, focussing on women since nursing and midwifery is a

---

[7] The then Financial Secretary to the Treasury in 1981, Nigel Lawson, later to become Chancellor of the Exchequer, described the 'massive public sector pay commitments' of the 'Clegg Commission' as an 'explosion', and the recommendations that the Commission made as "recommendations accordingly – which the Conservative Opposition, in the heat of the pre-election period, had pledged itself to honour." (Lawson, 1981)

[8] The latest report of the Review Body, which contains a fair amount of labour market information, is available as NOHPRB (2007).

female-dominated profession. Figure 4 illustrates the raw differential for female NHS nurses and midwives relative to the control group of women who have ever been a NHS employee but are currently employed elsewhere, with the 'events' described in the brief history above highlighted in the chart. As expected the raw differential falls during the period of the Social Contract, rises sharply as a result of the Clegg Commission reports, and then slowly declines until the jump associated with the regrading exercise of 1988, after which there is again a slow erosion of average public sector nurses' pay until the end of our data period.

We apply the method described in Section 2 to the panel data in order to assess the magnitude and trajectory of the nurses' pay differential. Our method tracks almost exactly the impact of the Social Contract and Clegg at the beginning of the period exhibited in the raw data, but thereafter, there are interesting divergences in both trend and cycle. Our method suggests that from Clegg until the 1988 regrading, the pay differential remained approximately zero, rather than the trend downwards in public nurses' pay observed in the raw data. After the 1988 regrading, our method predicts a positive premium for most of the remainder of the period, of just below 10% by the 'first stage' method and just above 10% from the 'second stage' method. This is in contrast to the raw data, which shows a steady decline in the nurses' 'premium', indeed becoming negative by the 2000s. This decline continues even after the introduction of 'Agenda for Change' in 2004-05 – our careful analysis of the data suggests that this decline arises because the earnings of the control group, particularly those not employed in health care, rose *faster* than the growth of pay of nurses in that last period.[9]

How do we interpret the relative stability in our predicted wage differentials in the 1980s, and after the 1988 regrading, compared to the apparent decline in the raw differential? Since our method controls for changing workforce *composition*, including changing composition explicitly arising from pay differentials, it seems likely that our measure is capturing a pay-workforce quality trade-off – akin to the Nickell-Quintini (2002) result suggesting that declining relative teachers' pay led to a fall in the average qualifications of new entrants to the profession. Indeed, this is confirmed by qualitative evidence: studies of the 1988 pay regrading exercise point to the difficulties in retaining and recruiting qualified nurses in the mid-1980s, just as the late 1990s era appear to

---

[9] Of course, analysis of the raw ASHE data without using these controls suggests an above-average rise in nurses' pay in 2004-05 – see NOHPRB (2007). Note that the ASHE data for 2006 do not include the 2.5% PRB pay award, which was paid (but backdated) some months later than these data were compiled.

have led to a greater use of non-qualified nursing staff and the failure of the NHS to retain higher paid (and presumably higher quality) more experienced nurses.[10] We therefore conclude that our estimation method is capturing the essential features of the evolution of the nursing workforce and its pay levels over the period.

4. Conclusion

This paper proposes a new method for analysing sector wage differentials using pay data. We utilise the standard fixed person effects assumption but, by using a two-stage estimator based around the residuals from testing the null hypothesis of a zero sectoral effect, we are able to construct time-varying estimates of these sectoral differentials. We also by-pass the tricky issue of constructing an identification strategy to control for worker selection (although the implications of selection are important in interpreting our results). This is because our method controls for changing unobserved as well as observed characteristics of the workforce. In principle, various extensions to our method are possible and the analytical properties of our test statistics also need further investigation.

Two important empirical results stem from our analysis. First long run public sector pay differentials do not seem to depart strongly from zero. Our second stage estimation procedure does *not* impose zero on the long run differential, but it seems implausible in any event that there would be large long run differentials in labour market rewards where in markets there is a high degree of worker mobility. Second, our estimation procedure captures 'episodes' such as pay regrading exercises, public pay policies, and the broad counter-cyclicality of the public-private wage differential. Moreover it also captures the change in composition of the public sector workforce that arises from disparate trends in public-private sector pay in periods where there is no dramatic upheaval in the pay structure. So, our finding that in the late 1990s, nurses' pay maintained a 10% premium relative to comparators should not be interpreted as

---

[10] The following description of the late 1990s is revealing, and consistent with the difference between the raw differential and our estimates of the pay differential: "The large proportion of hospitals' budgets that was spent on nursing, coupled with the disappearance of students as part of the labour force, led the NHS to attempt to control costs by substituting less skilled staff for registered nurses where possible. Enrolled nurses were offered conversion courses so that they could be registered and simultaneously the door was opened to the replacement of nurses by health care assistants, now with National Vocational Qualifications. A second-level nurse had provided basic nursing care for many years, and it was suggested that generic carers with comparatively brief training could provide most care in the future. A new, broader-based role encompassing nursing but not conforming to traditional job descriptions was proposed." (From Rivett (1998)).

saying that nurses are 'overpaid' in a simplistic market where worker quality is constant, but rather that the decline in nurses' pay relative to comparators induced a compensating deterioration in workforce quality. This suggests that public pay policy should not concentrate solely on pay levels, but also on what remuneration-quality 'mix' of the workforce best reflects the public interest.

References

Belman, D. and Heywood, J. (1989) 'Government wage differentials: A sample selection approach', *Applied Economics,* 21, 427-438.

Borland, J., Hirschberg, J. and Lye, J. (1998) 'Earnings of public and private sector employees in Australia: Is there a difference?' *Economic Record,* 74, March, 36-53.

Buchan, J. and Seccombe, J. (2005) *Past trends, future imperfect? A review of the UK nursing labour market in 2005 to 2005*, Royal College of Nursing: London

Disney, R. and Gosling. A. (1998) 'Does it pay to work in the public sector?' *Fiscal Studies*, 19, 4, 347-374.

Disney, R. and Gosling, A. (2003) 'A new method for estimating public sector pay premia: Evidence from Britain in the 1990s', CEPR *Discussion Paper* 3787, London.

Disney, R., Goodman, A. Gosling, A. and Trinder, C. (1998) *Public Pay in the 1990s,* Commentary No. 72, London: Institute for Fiscal Studies.

Dustmann, C. and Van Soest, A. (1998) 'Public and private sector wages of male workers in Germany', *European Economic Review,* 42, 1417-1441.

Haskel, J. and Szymanski, S. (1993) 'Privatisation, liberalisation, wages and employment: Theory and evidence', *Economica*, 60, 161-181.

Jakubson, G. (1991) 'Estimation and testing of the union wage effect using panel data', *Review of Economic Studies*, 58, 971-991.

Lawson, N. (1981) *Thatcherism in practice: A progress report,* text of a speech given to the Zurich Society of Economics, 14 January 1981, Conservative Policy Centre: London.

Monteiro, P. (2004) 'Using propensity matching estimators to evaluate the impact of privatisation on wages', NIPE *Working Paper* 1/2004, Universidado do Minho.

Nickell, S. and Quintini, C. (2002) 'The consequences of the decline in public sector pay in Britain: A little bit of evidence', *Economic Journal*, 112, February, F107-F118.

NOHPRB (2007) Review Body for Nurses and Other Health Professions: *Twenty-Second Report*, Cm 7029, The Stationery Office: London.

Postel-Vinay, F. and Turon, H. (2007) 'The public pay gap in Britain: small differences that (don't) matter?' *Economic Journal*, 117, October, 1460-1503.

Rees, H. and Shah, A. (1995) 'Public-private sector wage differentials in the UK', *Manchester School*, LXIII, March, 52-68.

Rivett, G. (1998) *From cradle to grave: Fifty years of the NHS*, King's Fund, London, available at: http://www.nhshistory.net/cvrivett.htm.

Stelcner, M., van der Gaag, J. and Vijverberg, W. (1989) 'A switching regression model of public-private sector wage differentials in Peru: 1985-86', *Journal of Human Resources*, 24, 545-559.

Terrell, K. (1993) 'Public-private wage differentials in Haiti: Do public servants earn a rent?' *Journal of Development Economics,* 42, December, 293-314.

Van der Gaag, J. and Vijverberg, W. (1988) 'A switching regression model for wage determinants in the public and private sectors of a developing country', *Review of Economics and Statistics,* 70, May, 244-252.

Van Ophem, H. (1993) 'A modified switching regression model for earnings differentials between the public and private sectors in the Netherlands', *Review of Economics and Statistics,* 75, 215-224.

Figure 1

Public sector mean pay relative to private sector meanpay:
'Raw' differences in hourly earnings 1975-2006

Source: *New Earnings Survey (NES) / Annual Survey of Hours and Earnings (ASHE)*

Figure 2

Fixed effects estimates of public sector pay effect for men:
alternative specifications

Figure 3

Panel A

Estimates of public sector pay effect over time: men

Panel B

Estimates of public sector pay effect over time: women

Figure 4

## Pay of nurses in the public sector
– relative to all workers who have ever been nurses

Source: *Calculated from NES/ASHE*