

Public-Private Wage Differentials in Ireland, 1994-2001

By

Gerry Boyle, Rory McElligott and Jim O'Leary*
National University of Ireland, Maynooth

Corresponding author: Professor Gerry Boyle,
gerry.boyle@may.ie

Abstract

Are public sector workers in Ireland paid more than private sector employees, when such differences in productivity-related personal attributes and job characteristics are controlled for? We estimate that in 2001 the premium enjoyed by public servants was about 13 per cent. We find that the premium, is significantly bigger for those near the bottom of the earnings distribution than for those near the top, was significantly bigger for women than men in the mid-1990s but not at the end of the 1990s, and does not vary significantly across different levels of educational attainment. We estimate the premium for 2001 to be not significantly different from that estimated for 1994 despite this period a period of exceptionally rapid output and employment growth, and correspondingly sharp tightening of labour market conditions in the Irish economy. The most remarkable difference between our results and those of other researchers for other countries relates to the absolute size of the premium. A number of possible explanations for this difference are discussed.

* An earlier version of this paper was presented at the Irish Economic Association Annual Conference in Belfast in April 2004. The authors would like to acknowledge the helpful comments and assistance of their colleagues Denis Conniffe and Donal O'Neill and participants at a departmental seminar at NUI Maynooth. Any errors are our own responsibility.

JEL codes: J31, J45

Key words: public, private, wage, differentials, Ireland

Public-Private Wage Differentials in Ireland, 1994-2001

1. Introduction

Public sector pay is a matter of considerable policy interest. At the macroeconomic level, the public sector pay bill is a large element of government expenditure, accounting for up to 70 per cent of general government consumption spending in some OECD countries (see Jorges (2001)). As such, public sector pay is an important driver of the overall tax burden.

At the microeconomic level, rates of pay in the public sector have an important bearing on the recruitment, retention and motivation of public servants and, by extension, on the quality of public services. At the same time, public sector pay rates may influence rates of pay in the private sector and in this way too affect the international competitiveness of the economy.

Reflecting these concerns, the design of public sector pay determination systems is an ongoing preoccupation of policy-makers. The challenge is to identify and implement mechanisms that are capable of delivering outcomes that meet both efficiency and equity criteria. Consequently, recent developments in this field, in Ireland and elsewhere, have sought to institute arrangements that link rates of pay in the public sector more directly to private sector rates.

In tandem with these developments, an international body of research has emerged which focuses on the operation of public sector labour markets. An important strand of that research has sought to quantify public-private sector wage differentials, to explore their pattern across the earnings distribution, educational attainment levels and occupations, and to analyse their movement over time.

This paper is part of that strand of research and, as far as methodology is concerned, draws its inspiration from it. To that extent, it is concerned with simply applying a well-tried set of econometric techniques to the analysis of public-private sector earnings differentials in Ireland, focusing on the period from 1994 to 2001. However, there are aspects of the Irish experience, specifically its economic performance during the 1990s and institutional features of its pay bargaining system, that make an extension of this research to the Irish case a matter of especial interest.

The paper is organised as follows. In Section 2 we describe the economic and institutional background against which public sector pay has evolved in Ireland over the past decade or so. In Section 3 we discuss some of the contributions to the recent international literature on public-private sector earnings differentials. In Section 4 we describe the dataset we use to analyse public-private sector wage differentials in Ireland. In Section 5 we report our main results, and in Section 6 we summarise our main findings and draw some conclusions.

2. The Irish Context

The Irish economy experienced rapid growth of output and employment during the 1990s. Between 1993 and 2001 real GDP doubled while GNP increased at an average annual rate of almost 8 per cent in real terms. Over the same period, total employment in the economy increased by a cumulative 45 per cent, or at an annual average rate of almost 5 per cent (Table 1).

The principal counterpart of this vigorous expansion of employment was an increase in the labour force which, thanks to a combination of strong natural increase, rising female participation rates and the emergence of net inward migration, rose by 3 per cent per annum on average between 1993 and 2001 (Table 1). Unemployment also declined sharply, with the unemployment rate falling from 15.7 per cent in 1993 to 3.6 per cent in early 2001.

Table 1: Labour force trends, Ireland, 1993 and 2001

(000s)	1993	2001
Total Employment	1,183	1,717
Unemployment	220	65
Labour Force	1,403	1,782
Unemployment Rate (%)	15.7	3.6

Source: Budgetary and Economic Statistics, Department of Finance, Ireland.

Some of the credit for Ireland's economic performance during the 1990s has been given to the so-called system of "social partnership" that has been an integral part of the country's institutional architecture throughout the period. This system brings together the government and organisations representing employees and employers in a policy-oriented dialogue, the purpose of which is to find common ground in relation to social and economic objectives and the means of attaining them. The main outcome of the system has been a series of three-year programmes setting out agreed policies across a range of issues¹.

At the core of each of these partnership programmes has been a centralised wage agreement, characterised by broadly similar pay increases across the public and private sectors. Government has participated in these agreements, not only in its capacity as employer, but also in its role as fiscal authority. This has allowed the wage agreements to incorporate commitments by government in relation to income tax reductions.

In simple terms, the logic behind the proposition that social partnership played a key role in the expansion of the Irish economy during the Celtic Tiger era runs as follows: partnership permitted tax reductions which facilitated wage moderation which in turn improved cost competitiveness and boosted output and employment. Of course, proponents of the social-partnership model claim other conduits of causation as well, such as the positive effects of partnership on the industrial relations climate and its contribution to the predictability of labour costs. The point

¹ These programmes and the periods to which they applied were as follows. The Programme for National Recovery (1988-1990); The Programme for Economic and Social Progress (1991-1993); The Programme for Competitiveness and Work (1994-1996); Partnership 2000 (1997-2000); and The Programme for Prosperity and Fairness (2000-2002). The current programme, Sustaining Progress, covers the period 2003-2005. All of these programmes are published by Government Publications, Stationery Office, Dublin.

is that, whatever the conduit and whatever the logic, the belief that social partnership is a valuable performance-enhancing institution is a very strong one in policy-making circles².

Undoubtedly, reductions in income tax, through their effect on labour supply, played a role in ensuring low rates of wage inflation in the Irish economy, at least in the early Celtic Tiger years. However, by the late 1990s, tightening labour market conditions had started to exert clear upward pressure on wage settlements. Increasingly, wage increases exceeded the norm set out in the centralised agreement. For example, *Partnership 2000* – the programme that covered the period from 1997 through 1999 – provided for average annual increases of 2.2 per cent in basic pay across the public and private sectors, but by 1999 actual increases in earnings across many areas of the economy were running at over 6 per cent. Amongst construction workers, wage inflation was running at double-digit rates by 1999.

Against this background a number of concerns emerged on the part of public sector employees and employers. One was that public sector workers were being left behind in pay terms by their private sector counterparts. This perception appears to have been based more on anecdotal evidence than on the available statistical evidence³. Nonetheless, it helped to create a sense of grievance amongst public sector workers, and also raised concerns about the capacity of organisations in the public service to recruit, retain and motivate staff.

² A fairly typical view is the following from the Deputy Prime Minister and Minister for Enterprise, Trade and Employment, Ms. Mary Harney T.D.: “Social partnership has been a major factor underpinning Ireland’s economic and social progress since the late 1980s. It has played a vital role in increasing employment ... It has helped deliver sustained and sustainable increases in our standards of living. Above all, it has contributed to our dynamism and our confidence that we can successfully tackle common problems.” (*Speech to Manpower Ireland Group European Business Leaders’ Breakfast*, Dublin, April 2002.)

³ Irish CSO data show that between 1996 and 2000, for example, average weekly earnings increased at an annual average rate of 4.9 per cent for public sector workers. The equivalent rates of increase for industrial workers and employees of banks, insurance companies and building societies were 5 per cent and 4.1 per cent respectively. All of these groups experienced far lower rates of increase than construction workers: between 1996 and 2000 skilled construction workers enjoyed average weekly earnings’ growth of over 10 per cent per annum.

At the same time there was growing disenchantment with the existing public sector pay determination system. This was a system that theoretically provided for most rates of pay to be benchmarked against the private sector, with relativities within the public sector supposedly playing a modest supporting role. In practice, however, the system had evolved to the point where there were comparatively few grades whose pay was determined by direct reference to private sector comparators, and a great number of grades where internal relativities were the main determinants⁴. The resulting system had become rigid and inflexible, militated against the discrete treatment of individual grades and occupations, and had an inherent tendency to generate wage-wage spirals.

Against this background, the Irish government, together with the social partners, established the Public Service Benchmarking Body (PSBB) in July 2000. The purpose of the PSBB was to examine pay and jobs in the public sector for those grades not covered by the Buckley Review Group⁵, and to make recommendations in relation to rates of pay based, according to the Body's terms of reference, on 'in-depth and comprehensive research and analysis of pay levels across the private sector'.

The background to the establishment of the PSBB is dealt with more fully in O'Leary (2002). For present purposes, one aspect of its mandate is especially deserving of note, namely the scope of the exercise. While it was not asked to examine every single public sector grade, the number of grades that did come directly within its remit covered 145,000 public sector employees or over 60 per cent of the total. There is no international precedent that we are aware of for a benchmarking exercise of this scale or scope.

The PSBB reported in June 2002. It recommended pay increases for every grade it reviewed, ranging from 3 per cent in the case of

⁴ Amongst the grades where remuneration was explicitly benchmarked to the private sector were those at the very top of the public-sector hierarchy. Pay rates for these grades have been determined for many years now by the *Review Body on Higher Remuneration in the Public Sector* (the "Buckley Review Group"). The most recent report of this group was published in 2000 (Government Publications, the Stationery Office, Dublin).

⁵ See previous footnote.

laboratory technicians and third-level lecturers to 25 per cent for certain grades of paramedic, and averaging just under 9 per cent overall. The cost of implementing its recommendations is estimated at €1.1bn annually, or about 3.5 per cent of annual net current government expenditure. Implementation is now in train and, subject to conditions relating to changes in work practices and the like being fulfilled, the PSBB's recommendations will have been fully delivered by early 2005.

A curious feature of the PSBB's report is that it furnished no specific justification for any of the pay increases it proposed. Instead, it provided a generalised rationale for its corpus of recommendations that echoed its terms of reference and cited a number of broad considerations. At no stage therefore did the PSBB indicate that its pay recommendations, either in general or in particular, arose because of a pay gap between the public and private sectors and the perceived need to bridge such a gap. In this respect the PSBB's approach differed markedly from that of the Buckley Review Group, whose 2000 report contained explicit and unambiguous statements to the effect that salaries for top public service posts were substantially out of line with those of comparable positions in the private sector.

Still, the impression that the PSBB's research showed that the generality of public sector grades were underpaid relative to their private sector counterparts, and that such research results formed part of the rationale for the pay recommendations made by the PSBB, has been propagated by influential figures, including public sector trade-union leaders⁶. It is not possible to test the validity of these impressions since the PSBB's research has not been published.

⁶ For example, Peter McLoone, General Secretary of the IMPACT trade union, responded to the PSBB's report by saying: "The outcome has vindicated IMPACT's view that public service pay fell behind during the economic boom" (see the IMPACT website at www.impact.ie) The Executive Committee of another public service union, the Association of Higher Civil and Public Servants, responded as follows: "The independent Benchmarking Body has clearly established that public sector remuneration is behind, and in some cases substantially behind, equivalent private sector remuneration and the fair rate for the job" (see the AHCPs website at www.ahcps.ie)

One objective of this paper, therefore, is to examine the evidence on the extent of the pay gap between public and private sector workers at the time the PSBB's research was carried out and how that gap evolved over the preceding years. Another objective is to situate the Irish evidence in the context of the international literature on this issue. To this end, we use a variety of econometric techniques to investigate the size, nature and dynamics of public-private sector pay differentials in Ireland during the 1994-2001 period, and to compare and contrast the Irish experience in this regard with the international experience.

3. Some recent international findings

To our knowledge, the kind of analysis we have undertaken has never before been carried out on Irish data. However, as noted above, a considerable volume of research along these lines has been undertaken for other countries. Good reviews of the literature and of the methodological issues that arise are contained in Gunderson (1998) and Gregory and Borland (1999). A brief summary of methodology and principal research findings follows.

Ultimately, the question that the bulk of this research attempts to answer is whether employees who are identical in all material respects would earn more by working in the public than the private sector. The generic method used to address this question involves comparing the earnings of individual workers, controlling for differences in productivity-related characteristics and job attributes, across the two sectors. This typically involves estimating regression equations with some measure of earnings as the dependent variable, and with a set of explanatory variables designed to capture productivity-related personal attributes (such as age, experience and education) and job characteristics (such as occupation, establishment size and nature of contract). These are sometimes referred to as Mincer equations, after Jacob Mincer the economist who pioneered the analysis of earnings.

The most basic application of this general approach is the single-equation method. Typically an equation of the following form is estimated, using pooled data for public and private sector workers:

$$W_i = \alpha + \sum_{i=1}^N \beta_i X_i + \delta PUB + \varepsilon_i \quad (1)$$

where, W_i are individual earnings (usually in log form); X_i are variables describing personal attributes and job characteristics; PUB is a public-sector dummy variable that takes a value of 1 if the individual is a public sector employee and 0 otherwise; ε_i is a random error term; and α , β_i , and δ are parameters. The parameter δ is usually referred to as the “public-sector premium” if its value is positive and the “public-sector penalty” if its value is negative. Typically, this equation is estimated separately for males and females, although it can be estimated for males and females together in which case the X vector will include a gender variable. This equation may be separately estimated for different age strata, education attainment levels, occupations and so on.

OLS regressions of this basic type differ from each other in terms of the precise combination of explanatory variables used. Variables measuring age, experience and education are included as a matter of course. However, there is some debate about the appropriateness of including occupational variables and variables reflecting establishment size. This debate helps to clarify the nature of the exercise in question.

In relation to firm size, Gregory and Borland (1999) say:

“Whether it is appropriate to include firm size as an explanatory variable depends on what job characteristics of a worker are regarded as fixed i.e. characteristics that would not change if the worker switched between public and private sectors. For example, if it is believed that a worker observed in a job at a large-size firm will always work in a large-size firm, then it is appropriate to include firm size as an explanatory variable. On the other hand, if it is believed that a worker switching sectors enters a

‘lottery’ where the probability of obtaining a job at a firm of a given size depends on the distribution of firm size in that sector, then firm size should not be included”.

Melly (2002) summarises the debate in relation to occupational variables as follows:

“It is appropriate to control for such occupational differences if they are indicative of what is necessary to do the job. However, if occupational titles are inflated in the public sector to justify high wages, then it is not appropriate to control for such occupational differences”.

In the light of the kind of considerations referred to by these authors, many researchers do not take account of establishment size or occupational variables when analysing public-private earnings differentials. It is important to note that the exclusion of such variables can materially influence the research results. We return to this issue in greater detail when presenting our results for Ireland.

OLS regressions are, of course, estimated at the mean. Accordingly, if the earnings variable is in log form, the results indicate that public sector employees, controlling for all the variables in X , are paid δ per cent more or less (depending on whether the estimated coefficient δ is positive or negative) than their private sector counterparts. In accord with convention, we label this coefficient the public-sector “premium” if its value is positive and the “penalty” if its value is negative.

However, if the distribution of earnings around the mean differs as between the two sectors, the OLS estimate of δ will convey an incomplete picture of the differential. There are in fact solid *a priori* and empirical reasons for believing that the earnings distribution is more compressed in the public than the private sector. Motivated by this perception, a number of researchers have used quantile regressions to analyse public-private sector pay differentials.

Quantile regressions can be used to generate estimates of the public sector premium at different points along the wage

distribution. Typically the points chosen are the 10th, 25th, 50th (median), 75th and 90th percentiles. Such estimates indicate whether there is a systematic variation in the public sector premium or penalty across the wage distribution.

Gregory and Borland (1999) provide a comprehensive review of the empirical literature. Surveying research based on the dummy-variable approach up to the mid-1990s, they note that a significant positive public-sector premium ranging from 3-11 per cent was found for a broad range of countries. Approaches based on Blinder-Oaxaca-type decomposition generally corroborated these findings, and indicated that public sector workers received higher returns to productivity-related characteristics than their private sector counterparts.

Other common threads through the research reviewed by Gregory and Borland are that (i) the public sector premium tends to be higher for females than for males, and (ii) the premium tends to be higher for public sector employees at the bottom of the earnings distribution than for public sector employees at the top. Quantile regression analysis by Poterba and Reuben (1994) for the US, and by Blackaby, Murphy and O'Leary (1997) for the UK, indicate that the size of the public sector premium is inversely related to a worker's position in the earnings distribution. A not uncommon finding of this strand of research is that, in the case of males, the public-sector premium moves from being substantially positive at the 10th percentile to a moderate penalty at the 90th.

An interesting feature of the research surveyed by Gregory and Borland relates to the cyclicity of earnings in the two sectors: the cyclical component in earnings is larger for the private than for the public sector. A clear implication of this is that the average public sector premium is likely to be lower at the peak of the cycle (or, more generally, after a period of rapid economic growth) than at the trough. One might expect to find strong evidence of this tendency in the Ireland of the 1990s, given the extraordinarily rapid growth experienced by the economy during this period.

A considerable volume of research on public-private sector pay differentials has been undertaken in the period that has elapsed

since Gregory and Borland's review. Disney and Gosling (1998) use the British Household Panel Survey to estimate the premium for the UK for the period 1991 to 1995. Their OLS estimates produced a male premium of 4 per cent and a female premium of 19 per cent. Perhaps their most intriguing results concerned the estimated effects of third-level education. Their OLS estimates yielded a penalty of over 9 per cent for third-level educated men but a premium of nearly 7 per cent for women.

Jurges (2001), using the German Socio-Economic Panel, reports an average penalty of 3 per cent for men over the period 1984 to 1996 and an average premium for women of 10 per cent. His quantile regression results reveal a familiar pattern with the returns for public-sector employees narrowing as one moves up the earnings distribution. His time series analysis echoes another conclusion of the earlier research reviewed by Gregory and Borland, namely a tendency for private sector pay to be more sensitive to the economic cycle.

Melly (2002) also uses micro data to analyse public-private sector pay differentials in Germany in 2000. He finds a public sector penalty for men of 8 per cent on average and an average 9 per cent premium for women. Quantile regression analysis revealed a male premium ranging from 3 per cent at the 10th percentile to a penalty of 15 per cent at the 90th. In the case of women, the premium ranges from 17 per cent at the 10th percentile to zero at the 90th. Melly's results indicate that the premium, for both men and women, declines with educational attainment, in other words, that the returns to education are higher in the private sector.

Research by Lucifora and Meurs (2004) looks at the public sector premium in three European countries – France, Italy and the UK – using 1998 micro data. They employ two different specifications in their basic dummy-variable model which they estimated using OLS and quantile regression. In one specification they control for personal attributes only (age, gender, education etc.); in the other they add job characteristics (occupation, part-time) and a regional variable. Not surprisingly, across all three jurisdictions, the estimated public-sector premium is lower in the latter

specification, reflecting the greater explanatory power of a larger set of right-hand side variables.

Their OLS estimates of the public-sector premium, using the broader specification, range from 4.9 per cent in the case of Italy to 6.4 per cent for the UK. Their quantile regression analysis produces a familiar pattern of results: in each country the premium declines monotonically from the bottom to the top of the earnings distribution, becoming negative (albeit statistically insignificant) in each case at the 90th percentile. Separate analysis for males and females produces another familiar result: the premium across all countries and all points of the earnings distribution is higher for women than for men.

4. Description of the Irish data

4.1 The sample

The data used in our analysis are drawn from the Irish component of the European Community Household Panel (ECHP), an EU-wide project, co-ordinated by Eurostat. The ECHP is designed to produce a fully harmonised dataset relating to the social, financial and labour-market circumstances of households, based on longitudinal surveys. As such, the ECHP provides harmonised cross-sectional data for each year in which the survey is conducted, as well as longitudinal data that permit analysis of changes over time. The latest year for which ECHP data were available at the time of writing was 2001; the first year was 1994. For the 1994 survey a total of 9,900 individuals were interviewed. By 2001, attrition meant that less than half of this original group was interviewed. Our sample comprises a subset of the ECHP sample (Table 2). Since our interest is in comparing pay in the public and private sectors⁷, we exclude those individuals who are not in the labour force and those who are unemployed. We also exclude those engaged in agriculture, forestry and fishing. Finally, we confine our analysis to employees; in other words, we exclude the self-employed. There are several reasons for so doing. The first

⁷ In classifying employees, the ECHP uses a broad definition of the public sector which includes civil servants, teachers, nurses, army personnel, Gardai, local authority workers, employees of semi-state bodies etc.

is a reason based on principle, namely that self-employment is a materially different form of economic status from being an employee, not least in terms of dimensions like risk and security of employment. The second reason is a more pragmatic one: the quality of the earnings data that are available for the self-employed is questionable. In any event, most international studies of public-private sector earnings differentials are confined to employees.

Table 2: The sample size used in the analysis of Irish public-private wage differentials, 1994-2001

	1994	1995	1996	1997	1998	1999	2000	2001
ECHP sample	9,904	8,531	7,488	6,868	6,324	5,451	4,529	4,023
Our sub-sample	3,246	2,582	2,225	2,099	2,173	1,888	1,673	1,494

4.2 The dependent variable

The ECHP dataset contains a considerable number of income variables. Many of these are too broad for our purposes since they include transfer payments and other forms of unearned income. Some of the earnings series are also too broad, since they include earnings from more than one job.

Our analysis requires a gross earnings variable that relates only to principal occupation. The ECHP dataset offers just one such variable: gross monthly earnings. However, it is possible to construct a gross hourly earnings variable using the monthly data and an average weekly hours worked variable also reported. While an argument can be made for such an hourly measure, our judgement is that the monthly measure is the more meaningful of the two since the notion of an hourly wage has no relevance for a great many employees, and certainly not for the majority of managerial, professional and technical workers.

Accordingly, we use gross monthly earnings in our core specification. However, we also test the impact of employing the hourly measure. Since average weekly hours worked are higher in the private sector, the gap in hourly earnings between the two sectors is wider than the gap in monthly earnings.

A shortcoming of our earnings variable, common to other studies in this type, is that it excludes important elements of remuneration such as pension entitlements, benefits-in-kind and performance-related pay. As far as the latter is concerned, there is a series in the ECHP dataset called ‘other earned income’, but it is a deficient measure of performance-related pay in that it includes lump-sum payments such as redundancy payments, back-dated pay increases and the like. There are no data in the ECHP relating to pension entitlements, so it is simply not possible to construct a measure of remuneration that incorporates the value of pensions. The same applies to benefits-in-kind.

Of course, there are other dimensions of a job that, while not part of the remuneration package, are valued by employees. An obvious one, especially relevant to an analysis of public-private sector pay differentials, is job security. In general, public sector employment is more secure than private sector employment: the concept of ‘a job for life’ is much more prevalent in the public than the private sector, and the risk of losing one’s job is much lower. There is no way of adjusting our dependent variable to reflect this difference.

4.3 The explanatory variables

All the independent variables that are typically used in this kind of analysis are readily available in some form or other from the ECHP dataset. Respondents are asked their age, gender and marital status. These are straightforward and require no elaboration. Respondents are also asked whether they suffer from any chronic disability and, since this is an obvious potential influence on productivity, we have used it as an independent variable.

The education variable in Mincerian equations is usually either a ‘completed years of education’ or a ‘highest level of education attained’ variable. The former is preferable on grounds of

granularity. However, the ECHP data set does not yield a satisfactory variable of this type, so we have been constrained to use an attainment variable that distinguishes between three levels of completed education: lower second-level; upper second-level; and third-level (degree or other third-level qualification). In reporting our results below, these are designated as Level 1, Level 2 and Level 3 respectively. Our education variable therefore, does not distinguish between different third-level qualifications and, since the span here is very wide (from certificates earned after a two-year course at an Institute of Technology to a Ph.D. earned after up to ten years study at a university), this is not entirely satisfactory.

A work-experience variable is also typical in Mincerian equations. Conceptually, there are several options here, one being the total years worked by an individual across all employments (a measure of total work experience); another being years worked by an individual in his/her current employment (a measure of job-specific experience). The Living in Ireland Survey (LIS), from which the ECHP is derived, yields data that allow a measure of the former type to be constructed. This is not the case with the ECHP data set. As a result, we are constrained to using a measure of the latter type, which is arguably inferior. However, given that we also use age as an explanatory variable, we do not think this poses a serious problem.

The variables discussed so far - age, gender, marital status, disability, education and experience – reflect individuals' productivity-related attributes. We have also used variables that relate to job characteristics: unit size; occupation; contract type and part-time status. As far as unit size is concerned, the ECHP contains information on the number of people working in the establishment where the individual respondent is currently employed. The variable is stratified into six size ranges from less than three up to 500-plus. The contract-type variable distinguishes between permanent, fixed-term and casual contracts.

As for occupation, we distinguish between the following categories: 'managerial' (including legislators and senior officials); 'professional'; 'technical'; 'clerical'; 'service'

(including sales workers); ‘craft’; ‘operatives’ (including assembly workers), and ‘elementary’ (including unskilled labourers).

4.4 Differences between the sectors

On average, the public servants in our sample are paid much more than the private-sector workers (Table 3). In terms of our chosen dependent variable, gross monthly earnings, public servants received €2778 on average in 2001, compared with an equivalent figure of €1905 for private sector workers, a differential of 46 per cent. In terms of the gross hourly earnings variable we have constructed, the differential is substantially larger, at almost 60 per cent. Average hourly earnings for public-sector workers are estimated at €18.23 in 2001, compared with €11.41 for private-sector workers⁸.

The corresponding differentials for earlier years were of a similar order of magnitude. For the 1994-2001 period, our data set suggests that public servants achieved an annual average increase in monthly earnings of 6 per cent, slightly more than the 5.8 per cent indicated for private sector employees. The equivalent rates of increase in average hourly earnings for the public and private sectors are 6.4 per cent and 6.6 per cent respectively. These figures suggest that over the medium to long term the ratio of raw earnings as between the two sectors in Ireland does not change very much, even if there are appreciable short-term fluctuations in that ratio. This echoes the findings of other researchers such as FitzGerald (2002) and Casey (2004).

Table 3: Features of the data employed in the analysis of Irish public-private wage differentials, 1994 and 2001

	1994	1994	2001	2001
Variable	Public	Private	Public	Private
Observations #:	1,116	2,130	413	1,223
Gross monthly earnings €	1,847.52	1,281.19	2,777.55	1,905.25
Gross hourly earnings €	11.83	7.30	18.23	11.41

⁸ These figures imply that the average hourly earnings of public sector employees were almost 39% higher than the average for all employees in 2001. This is somewhat higher than the estimate of 34% arrived at by Casey (2004).

Age Years	39.95	33.05	41.73	35.13
Married %	0.718	0.488	0.685	0.488
Male %	0.555	0.581	0.504	0.546
Part-time %	0.148	0.116	0.094	0.136
Disability %	0.080	0.075	0.056	0.067
Level 3 Education	0.342	0.177	0.433	0.196
Upper Level 2 Education	0.420	0.491	0.387	0.491
Lower Level 2 Education	0.237	0.333	0.179	0.313
Unit size 1-4 employees	NA	NA	0.061	0.121
Unit size 5-19 employees	NA	NA	0.194	0.262
Unit size 20-49 employees	NA	NA	0.206	0.218
Unit size 50-99 employees	NA	NA	0.133	0.117
Unit size 100-499 employees	NA	NA	0.247	0.196
Unit size 500+ employees	NA	NA	0.145	0.076
Tenure <=3 years	0.242	0.454	0.288	0.602
Tenure 4-8 years	0.134	0.270	0.153	0.179
Tenure 9-13 years	0.143	0.098	0.114	0.074
Tenure 14-18 years	0.002	0.000	0.087	0.051
Tenure 19+ years	0.479	0.178	0.322	0.080
Contract type :	NA	NA	0.891	0.866
Permanent				
Contract type: Fixed	NA	NA	0.061	0.037
Contract type: Casual	NA	NA	0.027	0.064
Contract type: Other	NA	NA	0.022	0.033
Occup. type: Managerial	0.037	0.066	0.068	0.093
Occup. type: Professional	0.323	0.083	0.327	0.069
Occup. type: Technical	0.116	0.088	0.143	0.098
Occup. type: Clerical	0.186	0.143	0.186	0.153
Occup. type: Service	0.118	0.174	0.126	0.196
Occup. type: Craft	0.058	0.173	0.031	0.146
Occup. type: Operative	0.060	0.163	0.051	0.169
Occup. type: Elementary	0.101	0.110	0.068	0.076

Source: ECHP, Ireland.

On average, public sector workers in Ireland are older; more highly educated and have spent longer in their current job than their private sector counterparts. They are also more likely to be married and more likely to be female. In these respects, the Irish pattern mirrors that of other countries where similar research has been carried out (see, for example, Lucifora and Meurs (2004) for data on France, Italy and the UK, and Melly (2002) for data on Germany).

There are several other relevant dimensions along which Irish public sector and private sector workers are systematically different. Public sector workers are more likely to be employed on permanent or fixed-term contracts. They tend to work in larger entities: 39 per cent of public servants work in entities (local units) that employ 100 or more people; the corresponding proportion for the private sector is 27 per cent.

Importantly, the occupational profiles of the two sectors are very different. Almost 33 per cent of all public sector workers are classified as ‘professional’, compared with less than 7 per cent of the private sector. Here the high proportion registered for the public sector reflects *inter alia* the classification of teachers. At the other end of the occupational spectrum, 32 per cent of private sector workers fall into either the ‘craft’ or ‘operator’ categories, compared with just 8 per cent for the public sector.

One other difference between the two sectors is worth noting at this stage. It relates to the distribution of earnings within each of the sectors. *A priori*, one would expect the distribution to be more compressed in the public than the private sector. This is what the international literature shows. It is also what our data set indicates for Ireland. The coefficient of variation in respect of average monthly earnings is consistently larger for the private than the public sector over the period 1994 through 2001 (Table 4).

Table 4: Coefficient of variation in public and private gross-monthly earnings, 1994-2001

	1994	1995	1996	1997	1998	1999	2000	2001
Public	0.5283	0.4512	0.4562	0.4769	0.4811	0.5072	0.5590	0.5389
Private	0.7126	0.7056	0.6477	0.6462	0.6339	0.6177	0.6164	0.5893

Source: ECHP, Ireland.

However, in our sample, the coefficient of variation for the private sector declines substantially between 1994 and 2001, while it increases, albeit marginally, for the public sector over this period. Taking the two sectors together, our data set indicates that the coefficient of variation was considerably lower in 2001 than in 1994. What all of this suggests is that the dispersion of employee

earnings declined in the non-agricultural economy during the Celtic Tiger period and that this was more than fully accounted for by the reduced dispersion of private-sector earnings.

Other researchers have uncovered similar evidence. For example, Barrett, FitzGerald and Nolan (2002) report a small reduction in earnings dispersion as measured by the ratio of the top to the bottom decile for the period 1994 to 1997, in marked contrast to the large increase in dispersion they found for the period 1987 to 1994. They attribute the 1994-1997 pattern to immigration, which they argue was concentrated amongst skilled workers and thus augmented the supply of skilled workers relative to the supply of unskilled workers in the Irish economy. It is likely that a similar explanation applies to the period from 1997 through 2001 as well.

5. The central results

5.1 OLS regression results all employees

We first present OLS regression results for 2001 (Table 5). These are typical of all years of our analysis. Two sets of results are presented; the first based on gross monthly earnings as the dependent variable, the second on gross hourly earnings. The coefficients on all the explanatory variables have the correct sign and, with few exceptions, are statistically significant. Earnings are positively correlated with age, education attainment, job experience and establishment size, and vary as expected across the occupational spectrum. Employees who are male and married, all other things equal, tend to earn more than others. Those on permanent contracts, all other things equal, earn more than those on fixed-term or casual contracts.

Table 5: OLS estimates of the wage equation, 2001

Variable	Monthly earnings		Hourly earnings	
	Coefficient	Standard Error*	Coefficient	Standard Error*
Public sector	0.1298	(0.0243)	0.1636	(0.0234)
Age	0.0209	(0.0061)	0.0210	(0.0059)
Age squared	-0.0002	(0.0001)	-0.0002	(0.0001)
Male	0.2383	(0.0215)	0.1371	(0.0199)
Married	0.0783	(0.0251)	0.0813	(0.0251)
Part-time	-0.6518	(0.0337)	-0.0823	(0.0305)
Disability	-0.0178	(0.0345)	-0.0626	(0.0372)
Education level 3	0.3279	(0.0319)	0.3219	(0.0305)
Education level 2	0.1220	(0.0239)	0.1156	(0.0217)
Unit: 5-19	0.0416	(0.0358)	0.0678	(0.0363)
Unit: 20-49	0.0769	(0.0364)	0.1182	(0.0368)
Unit: 50-99	0.1301	(0.0404)	0.1580	(0.0406)
Unit: 100-499	0.1294	(0.0379)	0.1463	(0.0376)
Unit: 500+	0.1761	(0.0435)	0.1904	(0.0439)
Tenure: 4-8	0.0915	(0.0247)	0.0836	(0.0233)
Tenure: 9-13	0.0934	(0.0363)	0.0515	(0.0357)
Tenure: 14-18	0.1555	(0.0374)	0.1334	(0.0379)
Tenure: 19+	0.2254	(0.0347)	0.1934	(0.0334)
Management	0.3854	(0.0483)	0.3019	(0.0453)
Professional	0.3772	(0.0484)	0.4638	(0.0455)
Technical	0.2361	(0.0441)	0.2258	(0.0401)
Clerical	0.1204	(0.0402)	0.1154	(0.0359)
Service	0.0153	(0.0409)	0.0177	(0.0349)
Craft	0.2379	(0.0437)	0.2175	(0.0394)
Operative	0.1637	(0.0409)	0.1037	(0.0374)
Permanent	0.1597	(0.0380)	0.0997	(0.0326)
Fixed-term	0.1419	(0.0695)	0.0460	(0.0646)
Constant	6.3307	(0.1162)	1.2853	(0.1118)
R²	0.6493		0.5536	

* Heteroskedastic robust standard errors. A complete set of results for the years 1994 to 2001 is available by request from the authors.

The estimate of the public-sector premium, using gross monthly wages as the dependent variable, is 13 per cent and highly significant. In other words, public sector workers in Ireland in 2001 were, on average, paid 13 per cent more than private sector workers, conditioning on a wide array of productivity-related job and personal characteristics. The estimate of the public sector premium using gross hourly earnings is somewhat higher, at over 16 per cent, and also highly significant. These estimates are considerably higher than those reported by Lucifora and Meurs for

France, Italy and the UK for 1998, by Melly for 2000 for Germany, and by Disney and Gosling for the UK for 1991-95 and 1990-99. This, despite the fact that these authors use specifications that are likely to explain a smaller fraction of the raw public-private earnings differential than ours.

Table 6 sets out the estimates of the premium obtained for each of the years 1994 through 2001 using different specifications of the wage equation. Each specification uses gross monthly earnings as the dependent variable but a different set of explanatory variables. This exercise illuminates two aspects of the estimated public sector premium: (i) its sensitivity to the inclusion of certain explanatory variables, and (ii) its behaviour over time. We first explore the question of the sensitivity of the results to changes in specification.

The specification labelled ‘Basic’ in the table contains the following explanatory variables: age, age squared, gender, marital status, education and experience, as well as dummy variables for part-time status and disability. This is the specification that is most typical of the one found in the international literature. To this basic equation we add serially the occupation variable, the contract-type variable, and the unit size variable.

Table 6: Estimate of the public-sector premium (all employees) for various specifications of the wage equation, 1994-2001

Equation type	1994	1995	1996	1997	1998	1999	2000	2001
Basic	0.1526 (0.0180)	0.1693 (0.0188)	0.1605 (0.0187)	0.1281 (0.0198)	0.1535 (0.0205)	0.0960 (0.0219)	0.1335 (0.0236)	0.1398 (0.0245)
Basic + Occupation	0.1407 (0.0174)	0.1509 (0.0184)	0.1471 (0.0186)	0.1215 (0.0194)	0.1407 (0.0205)	0.0905 (0.0211)	0.1200 (0.0231)	0.1398 (0.0250)
Basic + Contract Type	NA	0.1599 (0.0187)	0.1521 (0.0188)	0.1300 (0.0195)	0.1625 (0.0204)	0.1027 (0.0218)	0.1398 (0.0235)	0.1442 (0.0245)
Basic + Unit Size	NA	NA	NA	0.1049 (0.0192)	0.1325 (0.0199)	0.0769 (0.0215)	0.1106 (0.0232)	0.1257 (0.0240)
Extended*	0.1407	0.1460	0.1418	0.0995	0.1336	0.0831	0.1080	0.1298

(0.0174	(0.0183	(0.0185	(0.0183	(0.0199	(0.0207	(0.0221	(0.0243
))))))))

* The coefficients of the ‘extended’ equation are not comparable across the full period – see text for discussion.

Note: standard errors are in parentheses. Complete regression results are available on request from the authors.

The inclusion of the occupation variable has the effect of reducing the estimated public sector premium in each year except 2001, typically by about 1-1.5 per cent points. The inclusion of contract type has the effect of reducing the estimated premium for the years 1995 and 1996, but has the opposite effect in each subsequent year. The contract-type variable could not be included for either sector for 1994 because the relevant data were not collected for that year. The unit-size variable could only be included from 1997, since the relevant data were not collected for the public sector in either 1995 or 1996 and no data on unit size were collected for 1994. Once included, the unit-size variable has the effect of substantially reducing the estimated premium, typically by more than 1.5 per cent points.

Comparing the coefficients listed in the first and last rows of Table 6 from 1997 reveals that the overall effect of including all three variables, as we do in the ‘extended’ equation, is to reduce the estimated premium for each year. For instance, in 2001, our estimate of the premium would be 14 per cent rather than 13 per cent, had we omitted all three variables from our equation.

It is important to note that the coefficients estimated for the extended equation are not comparable across the full 1994-2001 period. This is because both the unit-size and contract-type variables are excluded from this specification for 1994 and the unit-size variable is excluded for 1995 and 1996. The coefficients estimated for the extended model are however comparable across the years 1997 through 2001. For this sub-period, there is no discernible trend in the premium, although it does fluctuate considerably from year to year. Of some interest, in light of the tightening of labour market conditions that occurred, is the fact that the 2001 premium is bigger than the 1997 premium, though not by a statistically significant margin.

To examine the behaviour of the public-sector premium over a longer span we have to use one of the specifications that is correspondingly time-consistent. Both the basic specification and the basic specification augmented by the occupational variable permit us to compare the premium over the full 1994-2001 period. In the former case there is some suggestion of a decline in the premium, but the difference between the estimates for 1994 and 2001 is not statistically significant. In the latter case, the premium estimated for 1994 is virtually identical to that estimated for 2001. Again, given the tightening of labour market conditions that occurred over this time frame, and the widespread belief that this had lowered earnings in the public sector relative to the private sector, this is a noteworthy result.

5.2 OLS regression results: interaction effects

A disadvantage of the simple dummy-variable approach is that estimated differences between the two sectors are confined to the intercept term; in other words that returns to productivity-related personal attributes and job characteristics are constrained to be the same across the two sectors. One way of mitigating this limitation is to run regressions where the public sector dummy is allowed interact with various explanatory variables. This can be done with any or all of the explanatory variables. Here, we report on the results of two such exercises where the public sector dummy variable is interacted with the gender and education variables respectively. Otherwise, the specification of the equations is exactly as above: in particular, these equations are based on the specification that includes the full set of explanatory variables and uses gross monthly earnings as the dependent variable (the extended specification).

The results in respect of gender are summarised in Table 7. The second row contains the estimated public sector premium for women. We include the estimated premium for all employees in the first row to facilitate the interpretation of the results. The third row shows the estimated difference between the premium for women and the premium for men. The premium for women is positive and highly statistically significant throughout the period under review. It is also larger than the premium estimated for all

employees together. The premium estimated for men is consistently smaller than that estimated for women throughout the period (this is indicated by the negative co-efficient on the male-female difference given in the table). However, and this is a very important qualification, the female co-efficient is in a statistical sense significantly different from the total population co-efficient only for the 1994-1997 sub-period and the male-female difference is only statistically significant for the years 1994 through 1998. In other words, for the most recent years, our analysis indicates that the premium enjoyed by female employees in the public sector is not significantly different from that enjoyed by their male counterparts.

These are striking results. The findings in respect of the earlier years are in line with the international experience; the findings for the later years are not. The overall results suggest that something happened to change relative gender premiums around the turn of the century. What this might have been is beyond the scope of this paper to identify conclusively. It has been speculated in the international literature that the existence of a bigger public sector premium for women than men may reflect the more punctilious observance of gender equality legislation by public than private sector employers. It may be, therefore, that the decline in the significance of the relative female premium that we have estimated for Ireland may reflect the increasingly conscientious application of equality legislation by private sector employers. Another possible explanation, at least as far as 2000 and 2001 are concerned, is the introduction of the minimum wage, to the extent that female employees in the private sector benefited more than other groups.

Table 7: Estimates of the public-sector premium by gender, 1994-2001

	1994	1995	1996	1997	1998	1999	2000	2001
All employees	0.1407 (0.0174)	0.1460 (0.0183)	0.1418 (0.0185)	0.0995 (0.0183)	0.1336 (0.0199)	0.0831 (0.0207)	0.1080 (0.0221)	0.1298 (0.0243)
Females	0.1950 (0.0247)	0.1989 (0.0264)	0.2065 (0.0256)	0.1417 (0.0246)	0.1646 (0.0256)	0.1114 (0.0266)	0.1341 (0.0287)	0.1448 (0.0298)
Male-Female difference	-0.0963 (0.0296)	-0.0926 (0.0309)	-0.1151 (0.0303)	-0.0751 (0.0302)	-0.0567 (0.0310)	-0.053 (0.0338)	-0.0511 (0.0368)	-0.0304 (0.0407)

Note: Standard errors are in parentheses. Complete regression results are available from the authors upon request.

Turning to education, the results are summarised in Table 8. Public servants with Level 1 education consistently achieved a substantial and statistically significant earnings premium *vis-à-vis* their private sector counterparts throughout the period. The estimated premiums for those with Level 2 education are somewhat different: in some years higher, in some years lower. Critically, however, the difference between the premiums estimated in respect of Level 1 and Level 2 education is not statistically significant in any year. The estimated premiums for those with third-level qualifications evince a different pattern again: higher than the Level 1 premium for every year except 1999, but by a statistically significant amount in only two years: 1994 and 1998. Another perspective on this analysis is provided by comparing the estimates of the premium for those with Level 1 education with the estimates for all employees: the difference between the respective coefficients is not statistically significant in any of the years 1994 through 2001.

Table 8: Estimates of the public-sector premium by educational attainment, 1994-2001

	1994	1995	1996	1997	1998	1999	2000	2001
All employees	0.1407 (0.0174)	0.1460 (0.0183)	0.1418 (0.0185)	0.0995 (0.0183)	0.1336 (0.0199)	0.0831 (0.0207)	0.1080 (0.0221)	0.1298 (0.0243)
Level 1	0.1005 (0.0281)	0.1462 (0.0321)	0.1237 (0.0315)	0.1044 (0.0313)	0.0956 (0.0318)	0.0915 (0.0355)	0.0907 (0.0404)	0.0967 (0.0464)
Level 2- Level 1 difference	-0.0166 (0.0350)	-0.0287 (0.0379)	0.0023 (0.0378)	-0.0215 (0.0373)	0.0167 (0.0387)	-0.0023 (0.0440)	0.0049 (0.0481)	0.0471 (0.0541)
Level 3- Level 1 difference	0.1896 (0.0414)	0.0488 (0.0449)	0.0627 (0.0452)	0.0157 (0.0452)	0.0875 (0.0452)	-0.0218 (0.0472)	0.0423 (0.0525)	0.0370 (0.0602)

Note: standard errors are in parentheses. Complete regression results are available on request from the authors.

Overall, therefore, our analysis suggests that the returns to education in Ireland are not materially different as between the private and public sectors. This contrasts with the findings reported in the international literature. Poterba and Reuben (1994),

Disney and Gosling (1998 and 2003)⁹ and Melly (2002) each report results indicating that the returns to education are significantly lower in the public than the private sector. This is an aspect of our findings that requires further investigation. It may be that a more refined measure of education attainment would yield results closer to the international norm in this respect.

5.3 Quantile regression results

OLS regressions are estimated at the mean. If, as already noted, public sector and private sector earnings are distributed differently around their respective means, OLS estimates of the public sector premium will tell an incomplete story. Our basic statistical analysis of the raw data does suggest that the two distributions are materially different. Specifically, the distribution of earnings is significantly more compressed in the public than the private sector, as is the case in all other jurisdictions where similar research has been conducted.

In Table 9 we present estimates of the premium that are generated by regressions at the 10th, 25th, 50th, 75th and 90th percentiles, for each of the years 1994 through 2001. The estimate of the public-sector premium is evidently not constant across the earnings distribution. In 2001, for instance, the public-sector premium is estimated at 17 per cent at the 10th percentile, falls to a range in the vicinity of 8-9 per cent at the 75th and 90th percentiles¹⁰. This pattern is broadly in line with the pattern that has been documented for other countries. It indicates that the margin by which public sector workers earn more than their private sector counterparts, again controlling for all productivity-related variables included in the model, is greatest for low-paid workers and is smallest for employees at the top of the wage distribution.

Does our analysis yield evidence of a public sector penalty at any point on the earnings distribution? Well, we have estimated

⁹ Disney and Gosling's 1998 paper reports a 3rd-level penalty for public-sector men but a premium for women. Their 2003 paper, however, using a more extended dataset and an IV estimator to counter potential selection bias, indicates a premium for third-level educated men and women of 13 per cent and 17 per cent respectively.

quantile regressions at the 95th percentile and these show that the public sector premium is not statistically significant at that point for any year except 1994. Carrying out quantile regressions at points of the distribution higher than this is problematic because the standard errors become increasingly large. We can only surmise that at a point of the earnings distribution above the 95th percentile significant public sector penalties exist, as documented by the Buckley Review Group.

Table 9: Estimates of the public-sector premium across the wage distribution, 1994-2001

	1994	1995	1996	1997	1998	1999	2000	2001
OLS	0.1407	0.1460	0.1418	0.0995	0.1336	0.0831	0.1080	0.1298
(Mean)	(0.0174)	(0.0183)	(0.0185)	(0.0183)	(0.0199)	(0.0207)	(0.0221)	(0.0243)
10% percentile	0.1740	0.2574	0.2293	0.1656	0.2092	0.1480	0.1320	0.1688
	(0.0376)	(0.0440)	(0.0322)	(0.0310)	(0.0361)	(0.0298)	(0.0351)	(0.0457)
25% percentile	0.1719	0.1705	0.1592	0.1584	0.1791	0.1387	0.1319	0.1474
	(0.0223)	(0.0269)	(0.0257)	(0.0256)	(0.0246)	(0.0260)	(0.0252)	(0.0322)
50% (Median)	0.1179	0.1366	0.1315	0.0966	0.1319	0.0724	0.0772	0.1164
	(0.0161)	(0.0223)	(0.0189)	(0.0202)	(0.0194)	(0.0221)	(0.0206)	(0.0311)
75% percentile	0.0930	0.0865	0.0931	0.0686	0.0948	0.0183	0.0895	0.0863
	(0.0182)	(0.0203)	(0.0206)	(0.0230)	(0.0226)	(0.0271)	(0.0286)	(0.0267)
90% percentile	0.0820	0.0494	0.0604	0.0161	0.0719	0.0039	0.0656	0.0941
	(0.0269)	(0.0340)	(0.0252)	(0.0307)	(0.0328)	(0.0380)	(0.0434)	(0.0451)

Note: standard errors are in parentheses. Complete regression results are available on request from the authors.

We have already noted that the degree of dispersion of public-sector earnings increased between 1994 and 2001, both in absolute terms and relative to the private sector. In other words, while the public sector wage distribution remained more compressed than the private sector distribution towards the end of this period, the margin by which it did so was narrower than in the early part of the period. Thus, we might expect to discover that the public-sector premium declined more precipitately as one travelled up the earnings' distribution in the early years of the period than in the later years. The estimates of the premium generated by the quantile regressions offer some support for this expectation.

6. Robustness checks – potential estimation biases

Estimates of wage differentials are potentially afflicted by a number of sources of bias, namely, sample selection bias conditional on observed right-hand-side explanatory variables; specification bias driven by unobservable variables relating to productivities, tastes and sector-specific factors; endogeneity bias; and bias arising from measurement error. In this section we explore the possible impact of these sources of bias on the estimate of the public-sector premium.

6.1 Sample selection bias

Valid estimates of public-private sector wage differentials require that comparisons be made on a like-for-like basis. Leaving aside the issue of potential unobservable variables that we consider below, we examine here whether, given the sample of observations and variables available to us, we can generate a valid set of comparisons between individuals employed in the public and private sectors. We address this question by utilising the method of “propensity scores” (see Conniffe, Gash and O’Connell, 2000).

This method comprises two steps. In the first step we estimate a probit model where the dependent variable is the public-sector dummy and the regressors are the set as employed in the previous section. The results of estimating the probit model for 2000 are given in Table 10. (Essentially similar findings are obtained for the other years in our analysis and are available from the authors upon request.) We observe positive and significant coefficients for age, third-level education, length of tenure in employment, professional occupational status and unit size. Therefore individuals that have a higher probability of being employed in the public sector (that is, those with a higher propensity score), tend to be older, better educated, of professional employment status, working in a larger unit and have been working in current job longer than persons with lower propensity scores. Lower scores are associated with occupations such as management, craft and operators. These results are consistent for all years. We obtain a *pseudo* R^2 of 0.261 in 2000 and this is similar each year.

Table 10: Probit estimates for public-sector status, Ireland, 2000

Variable	Coefficient	Standard Error
Age	0.0164	(0.0045)
Male	-0.0857	(0.0874)
Married	0.0047	(0.0967)
Part-time	-0.0104	(0.1296)
Disability	-0.0192	(0.1643)
Education level 3	0.3824	(0.1330)
Education level 2	0.1592	(0.1064)
Tenure: 4-8	0.4228	(0.1113)
Tenure: 9-13	0.5787	(0.1350)
Tenure: 14-18	1.0279	(0.1584)
Tenure: 19+	1.1860	(0.1282)
Unit: 5-19	0.0616	(0.1564)
Unit: 20-49	0.2359	(0.1601)
Unit: 50-99	0.2828	(0.1750)
Unit: 100-499	0.3978	(0.1573)
Unit: 500+	0.6656	(0.1754)
Management	-0.6240	(0.2009)
Professional	0.5423	(0.1846)
Technician	-0.0793	(0.1811)
Clerk	0.1435	(0.1686)
Service	-0.1012	(0.1677)
Craft	-0.9353	(0.2047)
Operator	-0.6029	(0.1806)
Permanent	-0.1508	(0.2316)
Fixed	0.4844	(0.2786)
Casual	0.0849	(0.2877)
Constant	-1.8852	(0.3381)
Log likelihood = -719.82443		
Pseudo R² = 0.2610		

Note: standard errors are heteroskedastic robust. Complete regression results are available from the authors upon request.

In the second step we rank the propensity scores and divide the observations into 10 equal strata, with the lowest 10% of the propensity scores in the first strata and the second 10% in the 2nd strata and so on. Rosenbaum and Rubin (1983) demonstrate that the distribution of the observable characteristics will be same within each strata. Therefore if one stratum is comprised of all private-sector workers then there are no similar public-sector workers to compare them with so these observations need to be omitted from the regression sample. As a working rule, we decided to omit a stratum of observations if it contained fewer than 20 individuals from either the public or private sectors. In the years 1997 to 2000, this rule resulted in the first three strata (30% of our observations) being omitted from our regression sample; in 1996 we omitted the first two strata; and in 1994 and 1995 we excluded the first strata. Our ‘extended’ wage specification was then re-run for all years on this truncated dataset and the resulting premium estimates are given in Table 11.

As is apparent the results from the truncated sample are very similar in all years and they are not significantly different in any year. This would imply that sample selection bias, conditional on observable right-hand-side variables, is not an issue with our data.

Table 11: Public-sector premium estimates for the ‘complete’ sample and the ‘truncated’ sample generated by the propensity score method, Ireland, 1994-2000

	1994	1995	1996	1997	1998	1999	2000
Complete sample	0.1407 (0.0170)	0.1460 (0.0187)	0.1418 (0.0187)	0.0995 (0.0187)	0.1336 (0.0196)	0.0831 (0.0209)	0.1080 (0.0227)
Truncated sample	0.1489 (0.0171)	0.1479 (0.0186)	0.1389 (0.0184)	0.0999 (0.0186)	0.1406 (0.0202)	0.0840 (0.0217)	0.1078 (0.0228)

Note: standard errors are in parentheses. Complete regression results are available from the authors upon request.

6.2 Specification, measurement error and endogeneity bias

To the extent that workers, due to unobservable talents or tastes, select themselves into particular sectors these factors will have an impact on the estimated public-sector premium. Thus it could be that part of the effect we're capturing in the estimated premium arises from an unobserved productivity or taste variable. Hence the OLS estimate of the premium would be an overestimate of the true premium. It is important therefore to assess the potential importance of this factor.

The panel feature of our data allows us to use individual-worker dummy variables as a proxy for possible unobservable factors. The application of the fixed-effects estimator to our extended wage model implies that the unobservable factors are unchanged over time and that the returns to these unobservable factors are similar for public and private sector workers.

Given these assumptions, the application of the fixed-effects estimator will yield an estimate of the premium that will be determined by those employees that move sector over the seven years of our panel. (A non-zero coefficient could also be generated by the presence of a differential wage growth between the public and private sectors over the seven-year period. This possibility is tested below.) This raises a number of issues that impinge on the assessment of the resulting estimate of the premium. It seems reasonable that the potential problems of both measurement error and endogeneity bias will tend to be exacerbated with the fixed-effects estimator. The presence of both of these biases will tend to bias the fixed-estimator towards zero which will operate to offset the upward bias due to omitted unobservable variables that afflict the cross-sectional regressions. But, as Disney and Gosling (2003) persuasively argue, there is no strong reason to suppose, as does Freeman (1984), that OLS will provide an upward bound and the fixed-effects estimator a lower bound estimate of the premium.

In the presence of both measurement error and endogeneity bias the appropriate estimator would be instrumental variables (IVs). As always obtaining IVs with the requisite properties is never easy. Fortunately the ECHPS contains some qualitative variables that could serve as potential IVs. One in particular looks promising, namely, workers' satisfaction levels in response to a

question on job security. Responses classified by public and private-sector employees for a number of years data are shown in Table 12. These data show that there is a strong correlation between “job security” satisfaction levels and public-sector status.

Table 12: Satisfaction levels with job security in the public and private sectors, Ireland, 1994-2000

Satisfaction level	1994		1996		1998		2000	
	%		%		%		%	
	Public	Private	Public	Private	Public	Private	Public	Private
1	8.46	10.91	3.5	5.82	3.83	3.76	1.55	2.18
2	4.53	8.62	2.43	6.49	2.19	5.28	1.81	4.05
3	6.2	15.27	4.71	11.73	2.19	9.85	3.1	6.34
4	9.45	17.67	7.45	18.89	9.49	16.41	6.98	17.36
5	20.18	21.92	22.95	30.12	19.34	29.94	20.16	32.95
6	51.18	25.63	58.97	26.96	62.96	34.75	66.41	37.11

Note: ECHPS, Ireland.

With this variable as an instrument we estimated the wage equation for the balanced panel using 2SLS. In the first stage we regressed the public-sector status dummy on the “job security” satisfaction variable and all of the other right-hand-side variables in the extended wage equation. We then generated the predicted values for the public-sector status dummy and used this variable instead of the original variable in the estimation of the wage equation and the public-sector premium.

Table 13 presents various estimates of the premium that were obtained using the panel dataset. We first compare our OLS estimates of the premium for the balanced panel with the average value of the estimate obtained for the seven years from the unbalanced panel. The results are broadly similar which gives us some confidence that panel attrition does not appear of itself to be associated with any noticeable biases.

The fixed-effects estimator yields a substantially lower premium than the OLS estimator. On the face of it, this suggests that the

OLS estimates are subject to a sizeable upward bias due to omitted variables. We estimated separate models for men and women which produced very similar estimates of the premium but the women's premium was not statistically significant (Annex 8). We also allowed the premium estimate to vary over time which revealed a pronounced downward trend over the period from 1994 to 2000.

Table 13: OLS, fixed-effects and fixed-effects and 2SLS estimates of the public-sector premium, Ireland, 1994-2000

Estimator	Premium estimate
OLS average annual (unbalanced panel)*	0.1218 (0.0195)
OLS balanced panel	0.1472 (0.0120)
Fixed-effects	0.0474 (0.0211)
Fixed-effects and 2SLS	0.1296 (0.0540)

* The simple average (1994-2000) of the premium estimates generated by the 'extended' specification of the wage equation in Table 6.

Note: standard errors are in parentheses. Complete regression results are available from the authors upon request.

The combination of the fixed-effects and the IV estimator reveals an estimate of the premium that is quite close to the OLS estimate. This suggests that the fixed-effects and OLS estimators are both subject to a substantial downward bias due to a possible combination of measurement error and endogeneity bias. Taking the fixed-effects and fixed-effects-2SLS estimators together suggests that the positive bias that probably arises due to the omission of important variables is more than offset by the negative biases that are potentially generated by measurement error and endogeneity. All in all therefore we feel that our results provide reasonable grounds for attesting to the robustness of the OLS estimates.

6. Summary and conclusions

Irish public sector employees are on average much better paid than their private sector counterparts. The ECHP dataset that we have used for the analysis presented in this paper indicates that the margin in the year 2001 was 46 per cent when measured in terms of average gross monthly earnings, the measure we consider to be the most appropriate of those available.

This differential is very large but of limited relevance. In particular, it does not necessarily mean that pay is higher in the public sector on a like-for-like basis. The reason is that public servants are typically better endowed than private sector employees with the kind of attributes that attract higher pay: they tend to be older, more experienced and better educated; they also tend to work in more highly-skilled jobs and in bigger establishments.

The key question addressed in our paper, therefore, is whether public sector workers in Ireland are paid more than private sector employees, when such differences in productivity-related personal attributes and job characteristics are controlled for, in other words whether a public-sector premium exists. The answer is clear and unambiguous. We estimate that in 2001, the latest year for which we have completed our analysis, the premium enjoyed by public servants was about 13 per cent, on the basis of gross monthly earnings.

We have analysed whether and how this premium varies with gender, with educational attainment, and across the earnings distribution. In summary, what we have found is that the public sector premium is significantly bigger for those near the bottom of the earnings distribution than for those near the top, was significantly bigger for women than men in the mid-1990s but not at the end of the 1990s, and does not vary significantly across different levels of educational attainment.

We have also analysed whether and how the public sector premium changed between 1994 and 2001. This was a period of exceptionally rapid output and employment growth, and correspondingly sharp tightening of labour market conditions in

the Irish economy. On the basis of international evidence, it might be expected that in such circumstances the intense competition amongst employers for a declining pool of available workers would have pushed up earnings in the private sector more quickly than in the public sector, resulting in a contraction of the public sector premium. Our research uncovers no compelling evidence that this was in fact the case: we estimate the premium for 2001 to be not significantly different from that estimated for 1994.

Some of our results are consistent with the research findings for other countries. In particular, our finding that the public sector premium declines with the level of earnings is in keeping with the international pattern. In contrast, our findings that female public servants no longer enjoy a significantly bigger premium than their male counterparts, and that the public sector premium for those that have completed third-level education is not significantly different from the premium enjoyed by early school-leavers, run counter to the international evidence.

However, the most remarkable difference between our results and those of other researchers relates to the absolute size of the premium. The public sector premium we have estimated for Ireland is high by international standards. At 13 per cent for 2001, it compares with corresponding estimates in the range 4-6 per cent for France, Italy and the UK, using 1998 data (our Irish estimate for 1998 is also around 13 per cent). This contrast is all the more notable given that we use a specification that would tend to produce a lower estimate of the premium than the specification used in most of the published international literature with which we are familiar. It is also worth pointing out that our estimate for Ireland predates the substantial boost to public-sector wages and salaries that flowed from the so-called 'benchmarking' process.

Our paper leaves a number of important questions unanswered. Probably the most important is: why is our estimated public sector premium for Ireland so large? One possibility is that our estimates are biased upwards because of flaws in our methodology. For example, we may have omitted important explanatory variables, or the sample we have used may be biased – an important possibility to check for given the attrition in the ECHP panel over the 1994-

2001 period. We have carried out a range of rigorous robustness checks to establish whether this might be the case, and we are reasonably confident that our estimates are not materially affected by such biases. That being the case, the question remains: why, on a like-for-like basis, are Irish public-sector employees paid so much more than their private-sector counterparts?

One possible explanation, a variant of the omitted variables argument, has to do with the deficiencies in our measure of earnings. Recall that the variable we use falls well short of being a comprehensive measure of remuneration, excluding as it does benefits-in-kind, performance-related pay and the value of pension entitlements. Of course, it also fails to take account of different rates of PRSI contribution (most public servants make lower PRSI contributions than their private sector counterparts) and the value of the greater security that typically attaches to public sector jobs. However, the problem with an explanation along these lines is that, if all of these factors were taken into account, it is likely that the estimated public sector premium would rise rather than fall, except perhaps in the upper reaches of the earnings distribution.

Another possible explanation is that certain public sector workers require compensation for the hazardous, stressful and/or unpleasant nature of the work they do. This type of argument is often advanced in the discussion of the pay of groups like prison officers, Gardai, teachers and nurses. However, it is not entirely persuasive for at least two reasons. First, hazardous, stressful or unpleasant work is not peculiar to the public sector. There are also a great many private sector employees who have to endure such working conditions. Second, unless there is reason to believe that the incidence of hazardous, stressful and unpleasant working conditions is greater in public service employment in Ireland than elsewhere, the argument does not explain why the public sector premium is so much higher in Ireland than in other jurisdictions.

The notion that a large premium for public sector employees can persist is an uncomfortable one for classically-trained economists, since it suggests that markets are not functioning properly. If markets were functioning properly, private sector workers would compete for better-paying public sector jobs to the point where the

public sector premium would disappear. Herein lies another possible explanation. The fact that a large premium persists may be indicative of market failures such as information deficiencies, skill mismatches and barriers to entry. The argument here is that some combination of these factors may effectively prevent private sector workers from competing for more lucrative public sector positions, except across a limited range of recruitment grades.

Even this explanation, despite its intellectual appeal, is incomplete however. It may shed light on why a public sector premium persists but, in itself, it does not explain why (i) the premium in Ireland is so much higher than elsewhere, or (ii) why the Irish premium was maintained at such a high level through a period characterised by the type of labour market conditions that might have been expected to cause it to decline. An obvious area to explore for answers to these questions is the wage bargaining framework that was in place during this period.

A full exploration of this area is outside the scope of this paper. That said, an obvious hypothesis suggests itself. It is that Ireland's social partnership model, and in particular the enormous value that politicians and policy-makers attached to it, conferred on Irish public sector trade unions greater bargaining power than they would otherwise have had (and greater bargaining power than their international counterparts typically enjoyed). According to this hypothesis, public sector trade unions were thus enabled to more effectively preserve barriers to entry, and more successfully pursue pay claims that matched the increases taking place across the private sector during the Celtic Tiger period.

References

- Barrett, A., FitzGerald, J., Nolan, B., 2002. "Earnings Inequality, Returns to Education and Immigration into Ireland", *Labour Economics* (November)
- Blackaby, D., Murphy, P., O'Leary, N., 1996. "Public-Private Sector Hourly Earnings Differentials in the UK: A Decile-Based Decomposition of the QLFS". *Economics DP 96-16, University of Wales, Swansea*.
- Casey, B., 2004. "An Economy-Wide Perspective on Earnings Data in Ireland". *ESRI Quarterly Economic Commentary* (Spring).
- Conniffe, D., Gash, V., O'Connell, P.J., 2000. "Evaluating State Programmes: "Natural Experiments" and Propensity Scores". *The Economic and Social Review*, Vol. 31, No. 4, October, pp. 283-308.
- Disney, R., Gosling, A., 2003. "A New Method for Estimating Public Sector Pay Premia: Evidence from Britain in the 1990's". *CEPR Discussion Paper No. 3787* (Feb.).
- Disney, R., Gosling, A., 1998. "Does it Pay to Work in the Public Sector?" *Fiscal Studies*, 19, 4, pp. 347-374.
- FitzGerald, J., 2002. "The Macro-Economic Implications of Changes in Public Sector Pay Rates", *ESRI Quarterly Economic Commentary*, pp. 43-58 (Winter), Dublin.
- Freeman, R., 1984. "Longitudinal Analysis of the Effects of Trade Unions". *Journal of Labour Economics*, 2, 1, pp. 1-26.
- Gunderson, M., 1998. "Government Compensation: Issues and Options". *CPRN Discussion Paper No. W/03* (July).
- Gregory, R.G., Borland, J., 1999 "Public Sector Labour Markets". In Ashenfelter, O., Card, D., eds., *Handbook of Labour Economics, Vol. 3c*, pp. 3573-3630.
- Jurges, H., 2001, "The Distribution of the German Public-Private Wage Gap". *Mimeo, Dept. of Economics, University of Dortmund, Germany*.
- Lucifora, C., Meurs, D., 2004, "The Public Sector Pay Gap in France, Great Britain and Italy". *IZA Discussion Paper No. 1041*.
- Melly, B., 2002, "Public-Private Sector Wage Differentials in Germany: Evidence from Quantile Regression". *Mimeo, SIAW, University of St. Gallen, Switzerland*.
- O'Leary, J., 2002. "Benchmarking the Benchmarkers". *ESRI Quarterly Economic Commentary*, pp. 77-91 (Winter), Dublin.

Poterba, J.M., Rueben, K.S., 1994. "The Distribution of Public Sector Wage Premia: New Evidence Using Quantile Regression Methods". *NBER Working Paper No. 4734* (May).

Public Sector Benchmarking Body, 2002. *Report of the Public Sector Benchmarking Body*, Government Publications, Stationary Office, Pn 11727, Dublin.

Rosenbaum, P.R., Rubin, D.B., 1983 "The Central Role of the Propensity Score in Observational Studies for Causal Effects". *Biometrika*, Vol. 70, No. 1 (April), pp. 41-55.