Cronfa - Swansea University Open Access Repository

_____

This is an author produced version of a paper published in:
British Journal of Industrial Relations

_____

Cronfa URL for this paper:
http://cronfa.swan.ac.uk/Record/cronfa50294

_____

Paper:

Murphy, P., Blackaby, D., O'Leary, N. & Staneva, A. (2019). Understanding What Has Been Happening to the Public Sector Pay Premium in Great Britain: A Distributional Approach Based on the Labour Force Survey. _British Journal of Industrial Relations_

http://dx.doi.org/10.1111/bjir.12474

_____

http://www.swansea.ac.uk/library/researchsupport/ris-support/

Understanding What Has Been Happening to the Public-Sector Pay Premium in Great Britain: A Distributional Approach Based on the Labour Force Survey

Philip Murphy (iD)*, David Blackaby* (iD)*,
Nigel O'Leary (iD) *and Anita Staneva* (iD)

Abstract

This article investigates what has been happening to the public-sector wage differential in Great Britain over the period 1994–2017. The evidence indicates that apart from men in the lower part of the pay distribution, the public-sector pay premium has declined for all public-sector workers. This decline has coincided with a decline in the overall pay gap, which is associated with changes in the composition of public- and private-sector workforces. As the relative pay disadvantage experienced by public-sector workers at the top of the pay distribution has worsened over time this must raise serious concerns about the ability of the public sector to recruit and retain the staff it needs to deliver public services.

1. Introduction

The important part played by public-sector pay in UK Government finances has meant that the pay of public-sector workers has attracted increased interest since the financial crisis, particularly as concerns about the size of the fiscal deficit has brought issues around the size of the public-sector pay bill to the forefront of government policy. As a result, the government has pursued a policy of wage restraint in the public sector, which has been a consistent feature of the political landscape in the UK since 2010. In fact, this policy of restraint is only now just beginning to weaken, as the Government announced in 2017 that it planned to scrap the 1 per cent public-sector pay cap, and agreed

Philip Murphy, David Blackaby and Nigel O'Leary are at the University of Swansea. Anita Staneva is at the University of Sydney.

a wage increase above the cap to more than 1 million public-sector workers in the summer of 2018. Nonetheless, the long and sustained period of wage restraint experienced by public-sector workers in the UK has naturally raised concerns not only about the effects wage restraint is likely to have on the living standards of public-sector workers but also on the ability of the public sector to attract and retain a high-quality workforce.

The use of wage restraint in the public is not a new phenomenon in the UK, even in recent times. For example, following a period of expansion in public-sector employment between 1998 and 2005, the then chancellor of the exchequer Gordon Brown wrote to all Pay Review Bodies in November 2005 urging them to limit recommended wage increases to 2 per cent or less. This call for restraint was reminiscent of earlier policies pursued in the 1990s and 1980s, where either increases in public-sector pay had to be financed from productivity and efficiency gains or a pay cap was imposed on certain groups of public-sector workers for a chronological listing of the various policies in place since 1980, see Pyper *et al.* (2018).

The pursuit of pay restraint in the public sector in the UK in the 1980s and 1990s also coincided with labour market reforms that changed the terms and conditions of employment for many public-sector workers. These changes sought not only to introduce greater flexibility into public-sector payment systems but also to increase the exposure of the sector to the forces of competition by the introduction of contracting out and compulsory competitive tendering (Bender and Elliott 1999; Elliott and Bender 1997). However, as the relative pay of many public-sector workers fell, as a result of both tighter financial controls and changes to the terms and conditions of employment in the public sector, some commentators began to raise concerns about the ability of the public sector to recruit and retain staff. Nickell and Quintini (2002), for example, noted that since the mid-1970s there had been a significant decline in the relative pay of most public-sector employees in the UK, and that this decline coincided with a reduction in the quality and ability of men entering teaching and general administration in the early 1990s compared to entrants in the late 1970s. Purely from a retention and recruitment standpoint, therefore, it is as important today, as it was then, that pay in the public sector is sufficiently attractive to attract and retain individuals with the skills and qualifications needed to deliver public services. For example, commenting on the potential effects of recent episodes of wage restraint in the public sector in the UK, Emmerson *et al.* (2016) stated: 'This could result in difficulties for public-sector employers trying to recruit, retain and motivate high quality workers, and raises the possibility of (further) industrial relation issues'. Similarly, following a decline in the relative pay of teachers, the School Teachers Pay Review Body (2016) noted that 'recruitment and retention pressures have become more acute, creating a challenging climate for schools'.

While recruitment and retention issues associated with changes in the relative pay of public-sector workers are an obvious concern, they are set against the popular perception that public-sector workers in the UK are well

paid. This popular view is often encouraged by media coverage that focuses attention on comparing the average pay of public- and private-sector workers (i.e. public/private-sector pay gap). However, comparisons made on this basis typically mask important composition differences between the two sectors, which can lead to significant differences between the average pay of public- and private-sector workers. The importance of compositional differences between the workforces of the public and private sector has long been recognized in the academic literature but their importance has often been ignored in the popular press.

Bender (1998) and Gregory and Borland (1999) both provide extensive surveys of the early empirical work devoted to measuring differences in public/private-sector pay across a wide range of countries. While Lausev (2014) provides a survey of more recent studies and compares findings on public/private pay differences in Eastern European transition economies with findings from developed market-based economies. These surveys provide two useful services. First, they provide a theoretical justification of why public-sector workers' pay can differ from private-sector workers' pay. For example, while pay determination in the private sector is likely to be profit constrained, pay in the public sector is more likely to be constrained by budgetary considerations that are determined within a political process. Second, they provide a summary of the work undertaken on measuring differences in public- and private-sector pay that account for differences in the composition of the two sectors. Most of the empirical evidence provided in these surveys is based on individual micro data, which allow investigators to control for characteristics that might be expected to affect individual earnings. Findings on the size of the public-sector pay premium are found to differ across countries, to vary over time and to depend on both the specification of the earnings equation used and the estimation method employed in the study. Nonetheless, some stylized facts do emerge for the UK, which are often shared by other market-based developed economies. For example, in the UK the public-sector pay premium tends to be higher for women than it is for men, and varies across the pay distribution: typically being higher at the bottom of the distribution than at the top, where the male pay premium is often negative.

Despite the importance attached to public-sector pay, evidence on what has been happening to the public-sector wage premium over time and across the distribution in the UK is limited. The main aim of this article, therefore, is to redress this balance by providing evidence on what has been happening to the public-sector pay premium over the period 1994–2017. The analysis is based on data from the Labour Force Survey (LFS), and estimates are provided separately for full-time male and female workers using a distribution-based decomposition method suggested by Fortin *et al.* (2010) and Firpo *et al.* (2018). The remaining sections of the article are set out as follows. Section 2 provides a summary, with a focus on the UK, of estimates of the public-sector pay premium; Section 3 outlines the methodology used in the analysis; Section 4 describes the data; Section 5 presents the results and discussion and Section 6 offers some conclusions.

2. UK estimates of the public-sector pay premium

Rees and Shah (1995) used data from the General Household Survey for 1983, 1985 and 1987 and estimated the public-sector pay premium on the basis of an Oaxaca (1973) and Blinder (1973) decomposition adjusted for sector selection (workers choosing to work in either the public or private sector). The Oaxaca–Blinder decomposition attributes any public/private-sector pay gap to either differences in the characteristics of workers employed in the two sectors or differences in the way characteristics are rewarded in the public and private sectors (the public-sector pay premium). They find that while characteristic differences for men largely explain the public/private pay gap, for women characteristics in the public sector tend to be rewarded more highly, which leads to a positive public-sector pay premium. Bender (2003) used data from the SCELI Survey for 1986 and found similar findings for estimates unadjusted for sector selection.[1] However, once sector selection effects were included the size of the public-sector wage premium increased substantially for both males and females.

Disney and Gosling (1998) estimated the public-sector wage premium using a simple intercept shift variable, indicating employment in either the public or private sector. Using data from the New Earnings Survey over the period 1979–1994, they find that once the occupation of workers is taken into account the public-sector pay premium fell for both men and women over the period as a whole, and that the premium had all but disappeared for men by 1994. Similarly, Postel-Vinay and Turon (2007), using data from the British Household Panel Survey between 1996 and 2003, find that while there was a small public-sector wage differential at specific points in time, the lifetime premium was 'essentially zero for highly employable individuals'.

More recently, while investigating the possibility of regional-based public-sector pay, Blackaby *et al.* (2018) estimated the public-sector pay premium by pooling LFS data for two time periods, 2009Q2–2011Q1 and 2011Q2–2015Q4. They find the pay premium is sensitive to the choice of the variables used in the earnings equation, and in particular on the inclusion of establishment size controls. Using the broadest possible specification of an earnings equation, they also find that over the two periods the public-sector pay premium fell for both male and female workers: from −3.9 to −4.1 per cent for men and from 5.6 to 2.4 per cent for women.

In the UK, and other Western developed economies, a common finding from studies that consider how the public-sector pay premium varies across the pay distribution is that while men in the public sector typically benefit from a positive pay premium at the lower end of the distribution, at the top a negative premium is more common. On the other hand, women employed in the public sector tend do much better than otherwise comparable women in the private sector at all points in the pay distribution (see Disney and Gosling (1998) and Blackaby *et al.* (1999) for the UK; Poterba and Reubon (1994) and Miller (2009) for the US and Mueller (1998) for Canada). Possible explanations for these findings include the success of trade unions in the

public sector in raising the wages of low-paid public-sector workers, political considerations that seek to limit the pay of more highly paid workers in the public sector and employers in the public sector pursuing less discriminatory pay practices. Irrespective of the exact cause, the tendency for employers in the public sector to pay above the going rate for low- and middle-skilled workers, and below the going rate for higher skilled groups is not new and was noted by Fogel and Lewin (1974).

Finally, Disney and Gosling (1998) find evidence that suggests the public–private-sector pay gap in the UK exhibits countercyclical behaviour: increasing sharply in the two recessions in the early 1980s and late 1980s and then decreasing as the economy moved towards a cyclical peak in the mid-1980s and 1990s. Maczulskij (2013) similarly finds evidence of countercyclical behaviour in the public-sector pay gap for Finland, suggesting that the greater sensitivity of private-sector wages to economic conditions may be a common feature of wage-setting behaviour in the public and private sectors.

3. Methodology

The analysis of pay gaps is often based on a decomposition framework, which decomposes differences in average pay into a composition and a pay structure effect. In these studies, the pay structure effect is typically seen as an estimate of the public-sector pay premium because it shows how characteristics are rewarded differently in the two sectors, while the composition effect reflects differences in the characteristics of workers in the two sectors. The decomposition of average wages into these two components is typically based on an Oaxaca (1973) and Blinder (1973) decomposition, which allows the composition effect to be further broken down into its constituent parts, showing the influence of particular characteristics. However, many of these decompositions were typically based on the sample means of the data, which is now seen as being too restrictive. As a result, several approaches have recently been developed that allow pay differences to be analysed at different points in the pay distribution, and depart from the simple intercept shift quantile regression method as used by Disney and Gosling (1998), which assumes characteristics are rewarded equally in the public and private sectors.

An excellent survey of distribution-based approaches to analysing pay differences is provided by Fortin *et al.* (2010). For example, DiNardo *et al.* (1996) use a non-parametric weighted kernel method to estimate counterfactual wage densities; while Machado and Mata (MM 2005) use a semi-parametric method based on conditional quantile estimates to do the same thing (see Melly (2005) and Lucifora and Meurs (2006) for an application of the MM method to analysing differences in public- and private-sector wages in Germany, France, Italy and Great Britain). Among the distribution-based approaches that are available, the use of a regression-based method proposed by Firpo *et al.* (2007 and 2018) has recently gained in popularity. Depalo *et al.* (2015), for example, use this method to estimate public/private-sector wage

differences in Euro-area countries. The approach of Firpo *et al.* can be briefly summarized as follows.

Let the difference in pay at the τth percentile (Q_τ) of the pay distributions of the public and private sectors — F_{w_g}, where $g = 1$ denotes the public sector and $g = 0$ the private sector — be given by:

$$D_{Q_\tau} = Q_\tau (F_{w_1}) - Q_\tau (F_{w_0}) \qquad (1)$$

The difference in pay between the public and private sectors at the τth percentile of the distribution can in turn be decomposed into two separate parts: (a) a part attributable to differences in the pay structure of the two sectors and (b) a part attributable to differences in the distribution of characteristics in the public and private sectors while keeping the pay structures the same.

In distribution terms, this is equivalent to a standard Oaxaca and Blinder decomposition.

$$D_{Q_\tau} = \left[Q_\tau (F_{w_1}) - Q_\tau \left(F_{w_0}^c \right) \right] + \left[Q_\tau \left(F_{w_0}^c \right) - Q_\tau (F_{w_0}) \right] = D_{S_{Q_\tau}} + D_{X_{Q_\tau}} (2)$$

where $F_{w_0}^c$ is the distribution of pay that would prevail if public-sector workers were paid according to the pay structure of the private sector. The first term in expression (2), therefore, is the part of the difference in public and private-sector pay at the τth percentile attributable to differences in pay structures between the two sectors ($D_{S_{Q_\tau}}$ the structure effect), while the second term is the part attributable to differences in the distribution of characteristics between the two sectors ($D_{X_{Q_\tau}}$ the composition effect). The counterfactual distribution and the other conditional distributions are found by applying relative weights to the data as described in Firpo *et al.* (2007 and 2018).

While the aggregate decomposition given by expression (2) provide insights into the overall role played by structure and composition effects in explaining pay differences at different points in the distribution, unlike the standard Oaxaca–Blinder decomposition it provides no information on the contribution different characteristics make to the overall composition effect. However, the method proposed by Firpo *et al.* (2007 and 2018) provides a solution to this problem by estimating the relationship between distributional statistics, based on recentred influence functions (RIFs), and a set of covariates using a standard regression model.

Coefficient estimates for the public sector, the private sector and the counterfactual distribution from the (RIF) regression model at each percentile ($\widehat{\gamma_{1,\tau}}$, $\widehat{\gamma_{0,\tau}}$ and $\widehat{\gamma_\tau^c}$) are then used to define the aggregate decomposition for the τth percentile as:

$$\widehat{D_{Q_\tau}} = \left[\overline{X_1}\widehat{\gamma_{1,\tau}} - \overline{X_0^c}\widehat{\gamma_\tau^c} \right] + \left[\overline{X_0^c}\widehat{\gamma_\tau^c} - \overline{X_0}\widehat{\gamma_{0,\tau}} \right] = \widehat{D_{S_{Q_\tau}}} + \widehat{D_{X_{Q_\tau}}} \qquad (3)$$

where a hat indicates that sample estimates have been substituted for population parameters, a bar indicates the use of sample means for the

public (1) and private (0) sectors and $\overline{X_0^c}$ is the sample mean of the reweighted characteristics of private-sector workers so that they have the same characteristics as public-sector workers. In practice, the difference between components of the decomposition shown in (3) and the equivalent elements in (2) are only of a second order of importance. Hence, the aggregate structure and composition effects reported in the analysis below are taken directly from the estimates given by (3).

Estimates of the pay structure $(\widehat{D_{S_{Q_\tau}}})$ and composition effects $(\widehat{D_{X_{Q_\tau}}})$ shown in (3) can be further broken down into a pure structure effect $(\widehat{D_{S,p_{Q_\tau}}})$ and a pure composition effects $(D_{X,p_{Q_\tau}})$, along with component that reflect either a specification error $(\widehat{D_{X,e_{Q_\tau}}})$ or a reweighting error $(\widehat{D_{S,e_{Q_\tau}}})$. In practice, both the specification and reweighting errors should be close to zero, and therefore act as a test of the model's performance.

For the structure and composition effects, we have:

$$\widehat{D_{S_{Q_\tau}}} = \overline{X_1}\left(\widehat{\gamma_{1,\tau}} - \widehat{\gamma_\tau^c}\right) + (\overline{X_1} - \overline{X_0^c})\widehat{\gamma_\tau^c} = \widehat{D_{S,p_{Q_\tau}}} + \widehat{D_{S,e_{Q_\tau}}} \tag{4}$$

$$\widehat{D_{X_{Q_\tau}}} = \left(\overline{X_0^c} - \overline{X_0}\right)\widehat{\gamma_{0,\tau}} + \overline{X_0^c}(\widehat{\gamma_\tau^c} - \widehat{\gamma_{0,\tau}}) = \widehat{D_{X,p_{Q_\tau}}} + \widehat{D_{X,e_{Q_\tau}}} \tag{5}$$

Given the linear nature of expression (5) the contribution of individual covariates $(j = 1, 2, \ldots, k)$ to the pure composition effect can then be shown to be equal to:

$$\widehat{D_{X,p_{Q_\tau}}} = \sum_{j=1}^{k}\left(\overline{X_{0,j}^c} - \overline{X_{0,j}}\right)\widehat{\gamma_{0,j,\tau}} \tag{6}$$

Of course, a similar breakdown of the pure structure effect is possible, but as in the standard Oaxaca–Blinder decomposition, where the model includes categorical variables, estimates of these individual components are sensitive to the choice of the omitted categories (see Jones 1983). For this reason, and because the composition effect is not affected by the omitted category problem, only a detailed breakdown of the composition effect is undertaken.

An important advantage of a linear decomposition, as shown in expressions (3)–(5), is that it can be used to highlight the part played by changes in both the structure and composition effects to changes in the overall pay gap at each percentile of the wage distribution over time. For changes in the wage gap that occurred between two time periods, therefore, we have:

$$\Delta_s\left(\widehat{D_{Q_\tau}}\right) = \Delta_s\left(\left[\overline{X_1}\widehat{\gamma_{1,\tau}} - \overline{X_0^c}\widehat{\gamma_\tau^c}\right]\right) + \Delta_s\left(\left[\overline{X_0^c}\widehat{\gamma_\tau^c} - \overline{X_0}\widehat{\gamma_{0,\tau}}\right]\right)$$

$$= \Delta_s\left(\widehat{D_{S_{Q_\tau}}}\right) + \Delta_s\left(\widehat{D_{X_{Q_\tau}}}\right) \tag{7}$$

$$\Delta_s\left(\widehat{D_{S_{Q_\tau}}}\right) = \Delta_s\left(\widehat{X_1}\left(\widehat{\gamma_{1,\tau}} - \widehat{\gamma_\tau^c}\right)\right) + \Delta_s\left(\left(\overline{X_1} - \overline{X_0^c}\right)\widehat{\gamma_\tau^c}\right)$$

$$= \Delta_s\left(\widehat{D_{S,p_{Q_\tau}}}\right) + \Delta_s\left(\widehat{D_{S,e_{Q_\tau}}}\right) \tag{8}$$

$$\Delta_s \left(\widehat{D_{X_{Q_\tau}}} \right) = \Delta_s \left(\left(\overline{X_0^c} - \overline{X_0} \right) \overline{\gamma_{0,\tau}} \right) + \Delta_s \left(\overline{X_0^c} \left(\widehat{\gamma_0^c} - \widehat{\gamma_{0,\tau}} \right) \right)$$

$$= \Delta_s \left(\widehat{D_{X,p_{Q_\tau}}} \right) + \Delta_s \left(\widehat{D_{X,e_{Q_\tau}}} \right) \qquad (9)$$

where $\Delta_s = x_t - x_{t-s}$ Similarly, the contribution of individual covariates to the change in the composition effect is:

$$\Delta_s \left(\widehat{D_{X,p_{Q_\tau}}} \right) = \sum_{j=1}^{k} \Delta_s \left(\left(\overline{X_{0,j}^c} - \overline{X_{0,j}} \right) \widehat{\gamma_{0,j,\tau}} \right) \qquad (10)$$

4. Data

The data used in the analysis are from the LFS. The LFS is a large-scale survey conducted by the Office for National Statistics (ONS), which has been undertaken on a quarterly basis since 1992. The current design of the LFS provides a sample of approximately 60,000 households in each quarter, and over the course of the survey respondents are interviewed on five separate occasions before being replaced by a new cohort. The rotating sample design of the LFS means that there is an approximate 80 per cent overlap of respondents from one-quarter to the next. To ensure individuals are only sampled once during their participation in the LFS, only those completing their fifth and final interview are included in the analysis. The fifth interview is also an occasion at which earnings-related questions are asked in the LFS.

  Information from quarterly LFS surveys is combined to produce annual data for the period 1994–2017. The year 1994 was chosen as the starting point of the analysis because a question on public-sector employment was not included in the LFS until the third quarter of 1993. Public-sector employment in the LFS is based on individual self-classification. Compared to employment returns from public-sector organizations reported in Public-Sector Employment (PSE) statistics, which is the preferred source of the ONS for estimating the size of the public sector in the UK, the LFS is known to overestimate the size of the public sector (Millard and Machin 2007). To examine this issue, Figure A1 in the appendix shows employment in the public and private sectors using the LFS and PSE. While Figure A1 suggests the LFS overestimates the size of the public sector, both series seem to track one another reasonably closely over most of the period analysed. Nonetheless, in order to consider the effect misclassification errors arising from self-classification may have on the results, a sensitivity test was undertaken that adjusted the LFS definition of public-sector employment used in the analysis. Specifically, following Millard and Machin (2007), individuals employed in either higher education or as agency workers were classified as private-sector employees regardless of how they were self-classified. Using this reclassification method had the effect of reducing the difference between LFS and PSE estimates of public-sector employment (see Figure A1).

Although the LFS is known to have weaknesses, its design and coverage makes it one of the best sources of data available in the UK to examine public/private-sector pay differences. In fact, the ONS regularly uses the LFS to investigate outcomes in the public/private sectors requiring information on pay and a wide range of worker characteristics. For this reason, and on the basis of the sensitivity tests reported later, we feel confident that it is the best source of data for the analysis undertaken in the article.

The measure of pay used in the analysis is the logarithm of real hourly earnings, which is found by deflating hourly earnings by the retail price index at constant 1990 prices. It is possible that the hourly earnings measure reported in the LFS may include an element of performance-related pay, for which no additional information is provided. As such payments are more likely to affect the pay of private-sector workers their influence might normally be expected to reduce the size of public/private-sector pay gap, particularly at the top of the pay distribution where performance pay can have a more significant impact on earnings. Of course, the size of any effect that performance-related pay has on the public/private-sector pay gap is likely to depend on underlying economic conditions, particularly as such payments tend to increase in an economic upturn and fall in a downturn. Interestingly, in a related study Bryson *et al.* (2017) find that the unexplained gap in the use of performance pay in the private and public sectors is significantly reduced when controlling for occupational status and other worker characteristics. This suggests that including occupations and other worker characteristics in the analysis might act as a useful control for unobserved performance pay effects.

While estimates of public/private-sector pay premiums reported in the literature often differ according to the specification of the earnings equation used, the econometric method employed and the country of study, there is nevertheless a broad consensus about the type of variables that should be included in the analysis. For example, most earnings equation estimates include demographic, human capital and workplace characteristics that have been shown to affect earnings. Consequently, in this analysis the construction of the counterfactual pay distributions and the RIF regressions used in the detailed decompositions included the following variables: age and age squared, age left full-time education, highest educational qualification, marital status, job tenure, occupation, region of work, establishment size and white ethnicity.[2] A definition for each of these variables is given in Table A1 in the appendix.

Both highest educational qualification and age left full-time education are included in the analysis to capture the influence schooling and certification effects can have on pay. Similarly age and age-squared are intended to capture the nonlinear nature of the age-earnings relationship suggested by a human capital model. Job tenure serves a similar role and can reflect both the acquisition of firm-specific skills and/or the presence of an internal labour market in the workplace, which is designed to promote co-operation and knowledge transfers between experienced and inexperienced workers

(Doeringer and Piore 1971; Topel 1991). Region of work, on the other hand, is included in the analysis to capture region-specific effects, which includes local labour market conditions that can affect pay.

While the decision to include most of the variables used in public/private-sector earnings equations is largely uncontested, others have attracted more critical attention. For example, Belman and Heywood (1989, 1990) have argued that establishment size should be included because public-sector establishments tend to be larger, and because pay and establishment size are positively related (Brown and Medoff 1989). Similarly, Moulton (1990) has argued that detailed occupation controls should be included to reflect the different distributions of occupation in the two sectors. However, others have argued that both establishment size and occupation should be excluded. Establishment size because it can represent the influence of monopsony power, and occupation because some occupations tend to be found almost exclusively in the public sector. Notwithstanding the rationale provided by Bryson *et al.* (2017) for using occupation as a control for the potential influence of unobserved performance-related pay effects, it seems likely that the broad definition of occupation used in the current analysis makes any problems relating to the specificity of public-sector occupations much less important. However, in order to investigate the sensitivity of the results to arguments relating to the use of establishment size and occupation two additional sensitivity tests were undertaken. First, establishment size and occupation were both excluded from the analysis; and second, public- and private-sector samples were matched on the basis of 3/4 digit occupation codes occurring in both sectors.[3] The results of these sensitivity tests are again reported later.

Of course, it is virtually impossible to control for every characteristic that might influence earnings. For example, it was not possible to include other job-related characteristics whose incidence and influence might vary by sector. This included job security, the risk of personal injury and fringe benefits, which includes pension entitlement. Elsewhere Disney *et al.* (2009) and Danzer and Dolton (2012) have used the concept of 'total reward' to calculate the present value of the remuneration package for different types of employees. For example, Danzer and Dolton (2012) compared total rewards for highly educated public- and private-sector workers and concluded that while rewards were broadly equalized for men, women had a clear advantage in the public sector. Elsewhere, Cribb and Emmerson (2014) have shown that when pensions and other benefits are taken into account the public/private-sector pay premium depends not only on what group of workers are being compared but also on what benefits are included. Given advocates of total reward analysis recognize the high demands of data needed to implement their approach, in line with much of the existing research on public/private-sector pay differences pension entitlement is not explicitly considered in the analysis. As a result, estimates of the public-sector pay premium reported below might underestimate the differential where public-sector pension provision is typically better than private-sector pension provision.

As pay structures are unlikely to be the same for full-time and part-time workers and are also known to differ for male and female workers, we restrict the analysis to full-time employees and present separate estimates for males and females. The means of the variables for 1994 and 2017 are available as an online appendix (Table OA). They show a fairly typical pattern for this kind of data: namely, that compared to private-sector workers, workers in the public sector tend to be older, better qualified, more likely to be in higher level occupations than in semi-routine or routine occupations, less likely to work in small establishments and have longer job tenure.

In addition to these workforce differences, Table OA also indicates that a number of demographic shifts have taken place in the data over time, which have affected the public and private sectors in similar ways. For example, the proportion of degree-holders (or more accurately degree and degree equivalent qualifications) and ethnic minorities has increased over time, while the proportion of full-time employees who are married (or living as married) has fallen for both males and females. Both features of full-time employment are consistent with changes that have occurred in the characteristics of the overall LFS sample between 1994 and 2017.

5. Results and discussion

All results reported in this section are based on an analysis of annual LFS data between 1994 and 2017. In each case, the analysis is undertaken separately for each year and separate estimates are also provided for male and female (full-time) employees. Results are reported at different points of the pay distribution, specifically at the 10th, 25th, 50th, 75th and 90th percentiles. All estimates are based on data using earnings weights provided in the LFS.

Aggregate Decomposition

Estimates of the aggregate decomposition given by (3) are shown in Figure 1 for males and females separately. At each percentile level, the graphs show the pay gap ($\widehat{D}_{Q_\tau}$), the part of the pay gap attributable to differences in pay structures between the public and private sectors ($\widehat{D}_{S_{Q_\tau}}$ the structure effect or public/private-sector pay premium) and the part of the pay gap attributable to differences in the distribution of characteristics of employees in the public and private sectors ($\widehat{D}_{X_{Q_\tau}}$ the composition effect). The graphs shown in Figure 1 report point estimates of the aggregate decomposition: additional information on the statistical significance of the individual components of the decomposition reported in Figure 1 is available from the authors upon request.

Some important features of the results reported in Figure 1 can be summarized as follows. First, for the period as a whole, the pay gap between public and private-sector workers has fallen across the pay distribution for both male and female workers. However, the fall has been higher at the top of the distribution than at the bottom, and also tends to be higher for female

FIGURE 1

Aggregate Decomposition Results by Gender and Percentile.
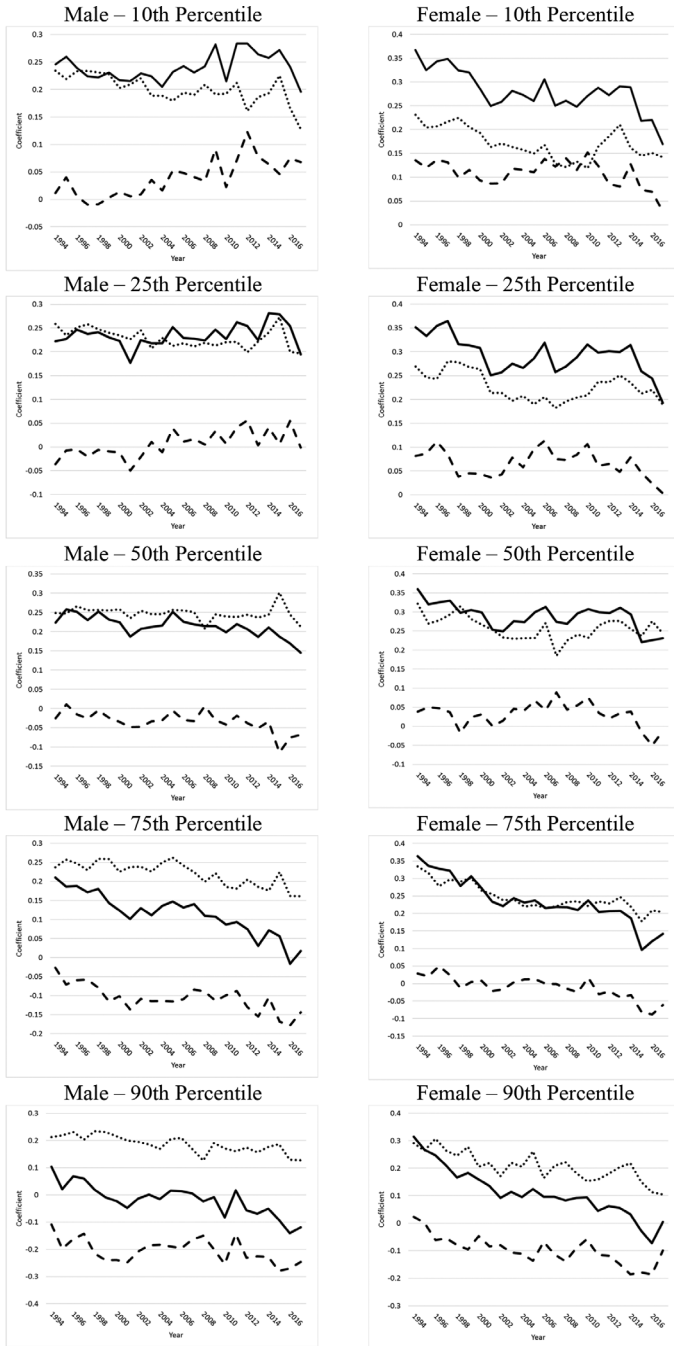
—— Overall difference   ······· Composition effect   – – Structure effect

workers than for male workers. Hence, for female (male) workers the overall pay gap fell by 31 (22) per cent at the 90th percentile between 1994 and 2017, compared to a 20 (5) per cent reduction at the 10th percentile over the same period.[4] The composition effect also falls for men and women over the period 1997–2017. However, while the decline in the composition effect tends to be larger at the top of the distribution than at the bottom for women, for men the decline follows a less distinctive pattern across the distribution, even though the largest falls are at the 10th and 90th percentiles. Differences in the experience of men and women are perhaps more interesting in terms of what has been happening to the structure effect (public-sector pay premium) over time. For example, while the public-sector pay premium increased between 1994 and 2017 for men at the 10th (a 6 percentage point increase from 1 to 7 per cent) and 25th (a 4 percentage point increase from −4 to −0 per cent) percentiles, it fell at every other reported percentile for both men and women. The largest reductions in the public-sector pay premium tended to occur at the top of the pay distribution, even though the reductions were almost as high at the bottom of the distribution than they were at the top for women. However, while positive public-sector pay premia are a common feature of the data at the bottom of pay distribution, at the top the premium is almost invariably negative for both male and female workers. Hence, at the top of the distribution not only do public-sector workers face a relative pay disadvantage compared to otherwise comparable private-sector workers but their position has also worsened over time. For example, at the 90th percentile, the public-sector pay premium for men fell from −11 to −25 per cent between 1994 and 2017, while for women it fell from 2 to −10 per cent over the same period.[5]

Second, when the pay of public-sector workers is compared to private-sector workers, female public-sector workers tend to fare better than their male counterparts. The relative advantage enjoyed by female public-sector workers is typically evident in both measures of the overall pay gap and the public-sector pay premium. However, as Figure 1 shows in the bottom half of the pay distribution — where a positive public-sector pay premium is more likely — the relative advantage enjoyed by female public-sector workers has fallen over the period as a whole: reflecting the markedly different changes in the pay premium enjoyed by men and women in the bottom half of the distribution, as noted in the previous paragraph.

Third, where there is a positive pay gap but a negative pay premium, the composition effect must by definition be larger than the observed pay gap. However, the change in the relative importance of composition and structure effects in explaining the overall public-sector pay gap over the period is perhaps more interesting. Thus, while the (absolute) magnitude of composition effects tend to be larger than structure effects at the bottom of the pay distribution, at the top of the distribution structure effects have become much more important over time. This is most evidently the case for men at the 90th percentile, where the (absolute) magnitude of the structure effect by the end of the period was significantly larger than the composition effect. An

increase in the relative importance given to structure effects is also evident at other points of the pay distribution. For example, the (absolute) size of structure and composition effects was almost equally important for women at the 90th percentile, and for men at the 75th percentile, by 2017. At the top of the distribution, therefore, differences in pay structures between the public and private sectors have widened despite government policies that have sought to make public-sector pay structures more like the private sector.

Fourth, and finally, while there have been reductions in both the public-sector pay gap and the public-sector pay premium over the period as a whole, these declines are neither constant from year-to-year nor do the changes observed always occur at the same time. However, some similarities in the profiles shown in Figure 1 are evident. For example, with the possible exception of men at the 10th and 25th percentiles, changes in the public-sector pay gap and in public-sector premium can be broadly characterized as follows: a period of decline between 1994 and the early 2000s; followed by episodes of improvement and decline between the early 2000s and the period immediately following the 2008/2009 recession; and then a further period of decline towards the end of the data period.[6]

These changes in the public-sector pay premium are most likely the result of a number of different influences, which can potentially pull in different directions at different times. However, the main contributors are likely to include (a) differences in workforce composition; and (b) differences in public/private-sector pay structures and how they respond to market conditions and government policy.

For example, it has been argued that the countercyclical behaviour often observed in public/private-sector pay is the result of private-sector pay being more responsive to market conditions than public-sector pay (Disney and Gosling 1998; Maczulskij 2013). However, the nature of this effect will depend on a number of other factors, which include the extent to which public-sector spending decisions directly influence the size of pay settlements in the public sector, the extent to which some groups in the public sector are favoured relative to others and the extent to which pay reform in the public sector has resulted in more flexible pay structures. The resulting situation is necessarily complex but it is nevertheless consistent with some of the changes observed in Figure 1. For example, the decline in the public-sector pay gap, and public-sector pay premia, which occurred between 1994 and the early 2000s, coincided with a period of economic recovery and a period of strict fiscal control. Fiscal conditions were tightened in the UK by the Conservative government in the early 1990s and continued until the end of the 1990s, following the decision made by the newly elected Labour government in 1997 to continue with the tight spending plans set out by the previous administration until 1999. Fiscal austerity, together with a curb on public-sector pay, either in the form of a public-sector pay freeze (2011–2013) or a 1 per cent public-sector pay cap (2013–2017), was also a distinctive feature of the period of decline witnessed in both the public-sector pay gap and public-sector pay premium towards the end of the period: when the economy was

slowly recovering from the effects of the 2008/2009 recession. Interestingly, in the much of the intervening period between these two episodes of decline in the public-sector pay gap and pay premium, the economy enjoyed a period of sustained economic growth and a commitment on the part of the government to invest in public services and the public sector. Hence, public-sector employment grew from the late 1990s until the mid-2000s, and even though Pay Review Bodies were urged by Gordon Brown in 2005 to keep pay settlements to less than 2 per cent, a formal policy of public-sector pay restraint and fiscal austerity was only introduced in the UK following the election of the coalition government in 2010.

The relative experience of public-sector workers shown in Figure 1, particularly in terms of seeing a positive pay premium at the bottom of the distribution and a negative pay premium at the top, is consistent with an argument that suggests pay setting arrangements in the public sector favour low-skilled workers over high-skilled workers (Fogel and Lewin 1974; Gregory 1990). In fact, extracts from official government policy announcements made in the UK during this period provide evidence to support the existence of such arrangements. For example, in 2009 the Labour government announced that it believed 'senior staff should show leadership in pay restraint' (Pyper *et al.* 2018). Similarly, when announcing the two-year public-sector pay freeze in 2010, the coalition government deliberately excluded public-sector workers earning £21,000 or less from the policy: instead these workers were to receive a pay increase of at least £250. Of course, it would be tempting to attribute the improvement in both the public-sector pay gap and pay premium at the bottom of the distribution to this policy. However, given that the improvement was seen only for men, and started prior to the introduction of the pay freeze, there is little to support this claim. Instead longer-term policy considerations, which have sought to favour lower-skilled public-sector workers, are more likely to have resulted in some public-sector workers having enjoyed a greater relative pay advantage than others.

Perhaps the strongest signal provided by the UK government that it wished to address inequalities in the pay distribution was the introduction of the National Minimum Wage (NMW) in April 1999. However, while improvements in the public-sector pay premium at the bottom end of the distribution occurred at around the same time as the introduction of the NMW, it seems unlikely that this could have been the cause of the changes reported in Figure 1. For example, the general consensus is that only about 4 per cent of workers are paid the NMW (Low Pay Commission 2012), and spillover effects have been found to be relatively small (Stewart 2012).[7] As a result, it seems unlikely that the NMW could have affected the percentiles of the pay distribution considered here; particularly as the analysis is based on full-time workers and the groups affected most by the NMW tended to be employed in the private sector, tended to be female and tended to be employed on a part-time basis (Bryson and Lucchino 2014).

The finding that women tend to do better than men, and that the public-sector premium is higher at the bottom of the pay distribution than at the top,

is consistent with much of the international evidence on public-sector pay in both the UK and in other developed economies. For example, differences in the estimated pay premium at different points of the pay distribution provide support for the argument that egalitarian and antidiscriminatory pay setting policies pursed in the public sector have resulted in the public sector having a more compressed wage distribution than the private sector (see Poterba and Rueben (1994) for the USA; Mueller (1998) for Canada; Disney and Gosling (1998) and Blackaby *et al.* (1999) for the UK and Lucifora and Meurs (2006) for an international comparison). In fact, some commentators have cited the push for greater equality in the public sector as the main reason why 'low wages' tend to be higher in the public sector and 'high wages' tend to be higher in the private sector (Yu *et al.* 2005). However, the finding that the relative advantage enjoyed by women in the public sector has fallen over time when compared to men, particularly at the bottom of the distribution, suggests a more complex process is at work. In particular, it seems that changes in pay structures at the bottom of the pay distribution have worked to improve the relative advantage of men in the public sector.

On the other hand, changes in the size of the negative public-sector pay premium at the top end of the pay distribution suggests that over the period as a whole differences in pay structures between the public and private sectors have resulted in an increase in the disadvantage faced by both male and female public-sector workers when they are compared to workers in the private sector. In fact, the decline in the public-sector pay premium seen at the top of the pay distribution represents a continuation of a much longer period decline in the UK: as identified by Disney and Gosling (1998) for the period 1979–1994 and by Yu *et al.* (2005) for the period 1994–2001.

Finally, it is worth noting that the movements in pay reported in Figure 1 are based on an analysis of the whole of the public sector, which is then compared to the whole of the private sector. However, as Dolton *et al.* (2015) point out public-sector employees are not a homogenous group, which is further complicated by the fact that their pay is determined by either the recommendations of Pay Review Bodies or by collective bargaining arrangements. Using a difference-in-difference approach Dolton *et al.* (2015) find evidence of year-to-year stability in the relative earnings of workers covered by different Pay Review Bodies over the period 1994–2007. While differences in the econometric approaches adopted make it difficult to directly compare the Dolton *et al.* (2015) estimates with our own, the finding that Pay Review Bodies have been unable in the short run to exert an independent influence on pay implies that the inclusion of appropriate occupation controls should allow different groups of workers in the public sector to be combined without any serious loss of generality.

Decomposition Based on RIF Estimates

Since RIF-based estimates of the reweighting error are typically small and statistically insignificant, estimates of the pure structure effect closely follow

the aggregate structure effect reported in Figure 1. As a result, estimates of the pure structure effect and reweighting error are not reproduced here. Instead, Figure 2 shows estimates of the composition effect as given by equation (5): the pure composition effect ($\widehat{D_{X,p_{Q_\tau}}}$) and the specification error ($\widehat{D_{X,e_{Q_\tau}}}$).[8]

Figure 2 indicates that while the pure composition effect tends to underestimate (overestimate) the composition effect at the bottom (top) of the distribution, both effects follow similar paths over the period as a whole, with the closeness of the two components being better at the top of the distribution than at the bottom. In fact, the relationship between the composition and pure composition effects are sufficiently similar to suggest that a detailed breakdown of the pure composition effect has value in identifying the factors that have contributed to changes in the composition effect over time.[9]

For the detailed decomposition of the pure composition effect (equation (6)), the variables included in the model are conveniently aggregated into eight separate groups: namely, age, education, job tenure, marital status, occupation, region of work, establishment size and white ethnicity. With two exceptions, the contributions made by these groups to the pure composition effect differ not only across the pay distribution but also over time. The two exceptions worth noting are: first, at the top of the distribution education consistently makes the single largest contribution to the pure composition effect over time; and second, occupation makes a significant contribution to the pure composition effect at a number of percentile levels over time. Because the time profiles of the other variables are more closely grouped over time, there is little to be gained in reporting year-by-year estimates of the detailed decomposition.[10] However, components of the detailed decomposition have more value in explaining changes that have occurred over time and it is to these changes that we now turn.

The falls in the public-sector pay gap that have occurred across the pay distribution between 1994 and 2017 is mirrored by falls in both the composition and pure composition effects (Figures 1 and 2). In order to examine these changes in more detail, Table 1 provides information on the changes that have occurred in the components of the pure composition effect between 1994 and 2017. For the sake of completeness, Table 1 also shows the changes that have occurred in the aggregate components of the decomposition between the same dates. The results shown for the aggregate decomposition are based on equation (7), and the result for the detailed decomposition of the pure composition effect on equation (10) using the eight groups of variables defined earlier. Table 1 is split into parts (a) and (b): part (a) shows the changes that have occurred in both the aggregate decomposition and the detailed components of the pure composition effect for males, while part (b) does the same for females.

With the exception of females at the 50th percentile, where the change in the pure composition effect does not reach a conventional level of significance, Table 1 indicates that all other composition-based changes — composition and pure composition effects — are significant at the 10 per cent level or better. Table 1 also highlights the components that have consistently contributed to

FIGURE 2
Breakdown of Composition Effect by Gender and Percentile.
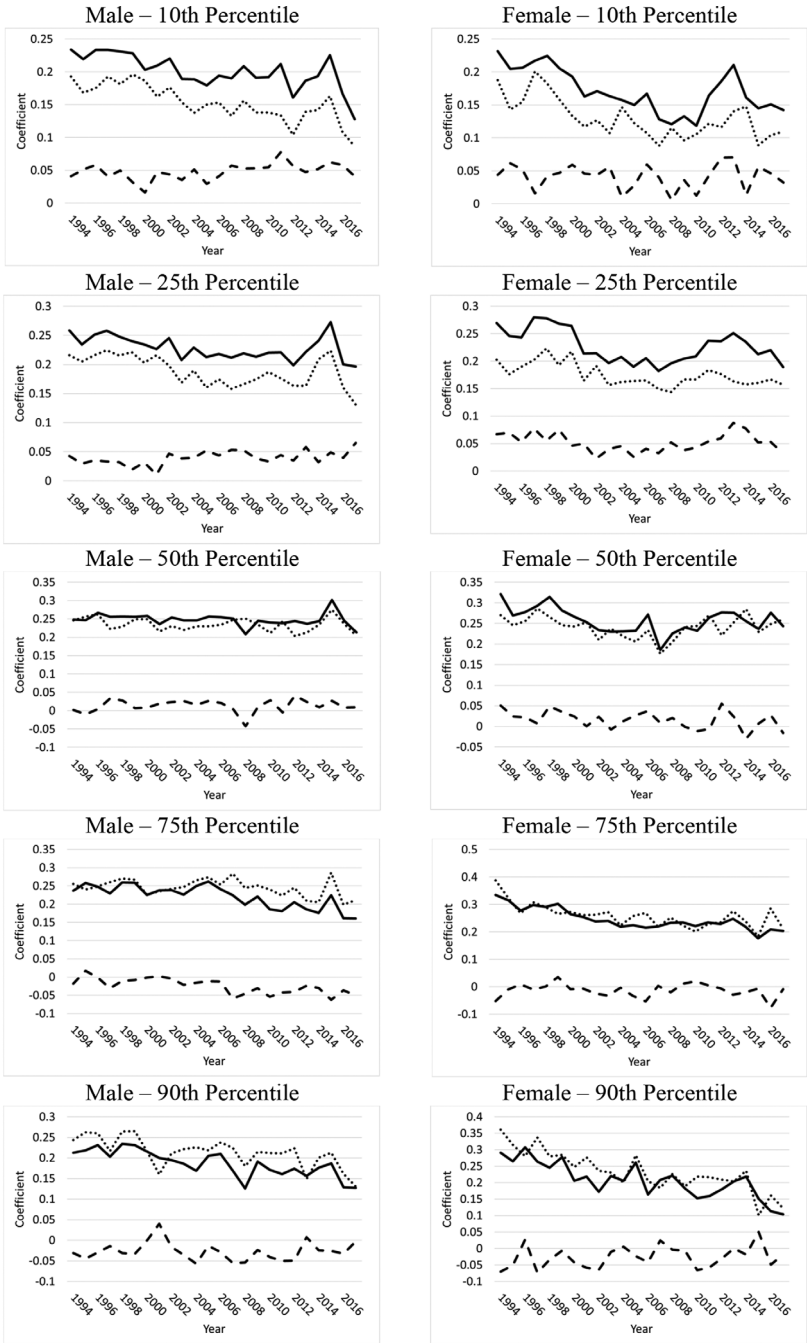
| —— Overall composition | ······ Pure composition | – – Specification error |

TABLE 1
Changes in the Aggregate and Detailed Decompositions by Percentile and Gender: 1994–2017

	10th	25th	50th	75th	90th
Part (a) Male					
Δ Overall	−0.0498***	−0.0277	−0.0782***	−0.1928***	−0.2232***
Δ Structure	0.0564**	0.0345	−0.0434	−0.1169***	−0.1383***
Δ Composition	−0.1062***	−0.0622***	−0.0347*	−0.0759***	−0.0849**
Δ Pure Comp	−0.1065***	−0.0850***	−0.0424***	−0.0441**	−0.1135***
Age	−0.0257***	−0.0095**	−0.0037	−0.0080*	−0.0122**
Education	−0.0029	−0.0079	−0.0034	−0.0248***	−0.0731***
Job tenure	−0.0284***	−0.0190***	−0.0121**	−0.0122*	−0.0217***
Marital status	−0.0104***	−0.0071***	−0.0004	0.0059**	0.0056*
Region of work	−0.0020	−0.0088***	−0.0214***	−0.0343***	−0.0319***
Occupation	−0.0185***	−0.0200***	−0.0032	0.0092	0.0080
Establishment size	−0.0198***	−0.0132***	0.0015	0.0195***	0.0105
White	0.0012	0.0005	0.0004	0.0006	0.0012
Part (b) Female					
Δ Overall	−0.1971***	−0.1585***	−0.1285***	−0.2212***	−0.3089***
Δ Structure	−0.1077***	−0.0781***	−0.0508*	−0.0894***	−0.1225**
Δ Composition	−0.0893***	−0.0804***	−0.0777***	−0.1319***	−0.1864**
Δ Pure Comp	−0.0778***	−0.0452***	−0.0106	−0.1759***	−0.2388***
Age	−0.0051	0.0001	0.0116	−0.0070	−0.0154*
Education	−0.0177	−0.0286***	−0.0437***	−0.1451***	−0.1996***
Job tenure	−0.0105	−0.0147**	−0.0021	−0.0044	0.0040*
Marital status	−0.0086	−0.0060	−0.0049	0.0036	0.0092
Region of work	−0.0021	−0.0030	−0.0173***	−0.0203***	−0.0177***
Occupation	−0.0231**	0.0112	0.0375***	−0.0276**	−0.0295**
Establishment size	−0.0110	−0.0046	0.0081	0.0222***	0.0075
White	0.0003	0.0004	0.0001	0.0027**	0.0026**

Note: Age (includes age and age-squared), education (includes age left school and highest educational qualification), tenure (includes work tenure variables), marital status (includes marital status variables), occupation (includes NS-SEC occupation codes), region (includes standard regions of work), establishment size (includes establishment size variables) and white. *, ** and *** denotes statistical significance at the 10, 5 and 1 per cent level, respectively.

a significant reduction in the pure composition effect. For example, in the bottom half of the pay distribution for men these were age, job tenure, marital status, occupation and establishment size. The equivalent components in the top half of the distribution were education, job tenure and region of work. For women on the other hand, few components were found to be consistently significant in the bottom half of the pay distribution, although occupation was significant at the 10th percentile and education and job tenure at the 25th percentile. By contrast, at the top of the female pay distribution, education, region of work and occupation all contributed significantly to a reduction in the pure composition effect over the period 1994–2017.

Table 1 indicates that changes in the education component were the largest single cause of the reduction in the pure composition effect for men and women at the top of the distribution over the period. For men, it accounted for 64 per cent of the change, and for women it was even higher at 84 per cent.

Together with the decline seen in the public-sector pay premium at the top of the distribution, these findings naturally reinforce concerns about the public sector's ability to attract and retain the high-quality staff it needs to deliver public services. In fact as Blackaby *et al.* (2015) have recently highlighted, the problems faced by the public sector in recruitment and retention are likely to have been exacerbated by changes in a number of 'non-pecuniary' aspects of public-sector jobs, which have reduced their attractiveness compared to the private sector. Changes in the relative attractiveness of public-sector jobs are serious and likely to have important implications for service delivery. For example, Propper and Van Reenen (2010) find that poor hospital performance in high earning areas in the UK was related to issues concerning the ability to recruit, retain and motivate NHS staff.

Sensitivity Analysis

To consider the sensitivity of the findings to changes in model specification and to misclassification errors arising from the self-reported measure of public-sector attachment used in the LFS, a number of sensitivity tests are reported in Table 2. The tests for the aggregate decomposition are based on three specifications: *specification 1* represents the baseline model, as reported in Figure 1, and is included for comparison purposes; *specification 2* follows Millward and Machin (2007) and reclassifies agency and higher education employees as working in the private sector; *specification 3* bases the analysis on a matched sample of 3/4-digit occupations in the public and private sectors and *specification 4* excludes both occupation and establishment size variables from the model. Two points from the results shown in Table 2 are worth highlighting. First, using the Millward and Machin (2005) reclassification method results in a reduction in both the overall pay gap and composition effect over time. Moreover, while some differences in the structure effect for specification 2 also arise (for example, men at the top of the distribution do slightly better and women slightly worse compared to specification 1), more generally changes in the components of the aggregate decomposition for this specification remain qualitatively similar to those reported for specification 1.[11] Second, excluding occupation and establishment size variables from the model has the effect of reducing the size of the composition effect and increasing the size of the pay premium in the bottom half of the pay distribution. The effects produced in the top half of the distribution are by comparison smaller and less significant. RIF estimates of the pure structure and composition effects for specifications 1 and 4 follow a similar pattern and indicate that while differences in these effects for the two models are statistically significant in the bottom half of the pay distribution, in the top half of the distribution they are not. Interestingly, RIF-based estimates for individual years between 1994 and 2017 also indicate that differences in specification and reweighting errors for the two models are not statistically significant in an overwhelming majority of cases.[12] Given there is little evidence to support the choice of specification 4 over specification 1 on the

TABLE 2
Sensitivity Test Results: Aggregate Decomposition

	10th			25th			50th			75th			90th		
	Over	Comp	Struc	Over	Comp	Struc	Over	Comp	Struc	Over	Comp	Struc	Over	Comp	Struc
Part (a)															
1994 (Males)															
Specification 1	0.2456	0.2340	0.0116	0.2225	0.2586	−0.0361	0.2236	0.2483	−0.0247	0.2102	0.2371	−0.0268	0.1043	0.2127	−0.1085
Specification 2	0.2352	0.2250	0.0102	0.2015	0.2394	−0.0379	0.1942	0.2206	−0.0264	0.1845	0.2042	−0.0197	0.0883	0.1580	−0.0697
Specification 3	0.2439	0.2325	0.0114	0.2166	0.2502	−0.0335	0.2173	0.2414	−0.0242	0.1989	0.2247	−0.0258	0.0922	0.2035	−0.1113
Specification 4	0.2456	0.1596	0.0860	0.2225	0.1922	0.0302	0.2236	0.2212	0.0024	0.2102	0.2413	−0.0310	0.1043	0.2213	−0.1170
2017 (Males)															
Specification 1	0.1958	0.1278	0.0680	0.1948	0.1964	−0.0017	0.1454	0.2136	−0.0681	0.0175	0.1612	−0.1438	−0.1189	0.1278	−0.2467
Specification 2	0.1882	0.1127	0.0755	0.1699	0.1597	0.0102	0.1025	0.1751	−0.0726	−0.0229	0.1068	−0.1296	−0.1506	0.0586	−0.2092
Specification 3	0.1984	0.1297	0.0688	0.1918	0.1964	−0.0046	0.1344	0.2066	−0.0723	0.0013	0.1442	−0.1429	−0.1313	0.1208	−0.2522
Specification 4	0.1958	0.0706	0.1252	0.1948	0.1297	0.0651	0.1454	0.1814	−0.0359	0.0175	0.1435	−0.1261	−0.1189	0.1162	−0.2351
Part (b)															
1994 (Females)															
Specification 1	0.3668	0.2314	0.1353	0.3513	0.2696	0.0817	0.3598	0.3215	0.0383	0.3636	0.3350	0.0287	0.3141	0.2907	0.0235
Specification 2	0.3568	0.2145	0.1423	0.3372	0.2534	0.0839	0.3321	0.3087	0.0235	0.3278	0.3169	0.0110	0.2648	0.2978	−0.0330
Specification 3	0.3673	0.2311	0.1362	0.3507	0.2695	0.0812	0.3558	0.3175	0.0383	0.3615	0.3357	0.0258	0.3109	0.2903	0.0206
Specification 4	0.3668	0.1504	0.2163	0.3513	0.1996	0.1517	0.3598	0.2909	0.0689	0.3636	0.3208	0.0428	0.3141	0.2942	0.0200
2017 (Females)															
Specification 1	0.1697	0.1421	0.0276	0.1927	0.1892	0.0036	0.2312	0.2438	−0.0126	0.1424	0.2031	−0.0607	0.0052	0.1043	−0.0991
Specification 2	0.1464	0.1287	0.0177	0.1685	0.1746	−0.0061	0.1883	0.2141	−0.0258	0.1041	0.1727	−0.0685	−0.0224	0.0732	−0.0956
Specification 3	0.1666	0.1427	0.0239	0.19206	0.1939	−0.0018	0.2275	0.2421	−0.014	0.1385	0.2038	−0.0653	0.0026	0.1001	−0.0975
Specification 4	0.1697	0.0609	0.1088	0.1927	0.1130	0.0797	0.2312	0.1796	0.0516	0.1424	0.1812	−0.0388	0.0052	0.1188	−0.1136

Note: Specification 1 is the full baseline specification; specification 2 follows Millward and Machin (2005) and classifies agency and higher-education employees as working in the private sector; specification 3 bases the analysis on matched 3/4-digit occupations in the public and private sectors and specification 4 excludes occupation and establishment size controls from the model.

basis of model performance, and given all the remaining sensitivity tests reported in Table 2 produce qualitatively similar results, our preference is for a model that includes the widest possible range of variables known to affect earnings: namely, specification 1. More generally, all the sensitivity tests reported in Table 2 provide strong support for the argument that the direction of change identified by the analysis is correct and not merely an artefact of misclassification errors or model specification.

6. Conclusion

This article has examined what has been happening to the pay of public-sector workers over the period 1994–2017. The analysis confirms that over the period, with the exception of men at the lower quartile and below, there has been a decline in the public-sector pay premium, which represents the continuation of an earlier trend that has seen the relative position of public-sector workers fall over time. Moreover, the position of public-sector workers is undoubtedly worse at the top end of the pay distribution where a negative public-sector pay premium is not only more commonly found but the pay disadvantage faced by public-sector workers has also grown over time.

While the public-sector pay premium has secularly fallen over time at most points in the pay distribution, there are still year-to-year variations in the premium. However, many of these changes seem to be broadly consistent with either changing economic and fiscal conditions or a policy pursued by wage setters in the public sector to favour low-skilled over more highly skilled workers. Thus, the relative pay of public-sector workers tends to move countercyclically, and low-paid public workers tend to do better than high-paid public-sector workers when compared to similar workers in the private sector.

Differences in the composition of the public sector and private sectors have made a significant contribution to the overall public-sector pay gap over time. For example, education and occupation are two of the most important compositional differences accounting for differences in pay between the two sectors: education at the top of the pay distribution and occupation across the distribution. However, other differences such as age, job tenure, establishment size and region have also made significant contributions at different points in the pay distribution, although their effect not only tends to be smaller but also their importance can vary over time.

Looking across the period as a whole, there has been a significant reduction in the effect that compositional differences have had on the relative pay of public- and private-sector workers. For men, a range of characteristics was identified, although the importance of particular characteristics depended on which part of the distribution was being considered. The picture for women, on the other hand, was less clear-cut, at least in terms of the statistical significance of different characteristics being able to explain how reductions in the pure composition effect affected the relative pay of public and private

sectors over the period as a whole. What does not seem to be in doubt, however, is that changes in the educational endowments of the public and private sectors was an important driver of changes in the relative pay of male and female public-sector workers at the top of the pay distribution. Of course, such changes might be expected over a period of almost 25 years, but they also coincided with a substantial decline in the public-sector pay premium, which significantly increased the pay disadvantage faced by public-sector workers at the top of the pay distribution. This raises serious concerns about the ability of the public sector to retain and recruit the high-quality staff it needs to deliver public services. The government must recognize, therefore, that if it wants to provide high-quality and sustainable public services it needs to adopt a public-sector pay policy that does not focus solely on the short-term goal of reducing pressure on the public finances. As noted by Runge *et al.* (2017), this should reduce the need for *ad hoc* responses to recruitment and retention difficulties caused by local and national labour shortages, and allow the public sector to develop a more coherent and strategic response to workforce issues.

Final version accepted on 2 May 2019

Notes

1. The SCELI Survey was based on a sample of six local labour markets that were chosen to provide contrasting patterns of employment in Great Britain.
2. In addition, the number of dependent children and interactions between age and age left school were included in the probit model for sector choice (public vs private sector employment) used to generate the counterfactual pay distributions. The decision to include the number of dependent children in the sector attachment equation reflects the common perception that public sector jobs offer more family-friendly working arrangements, which are more attractive to workers with children. Similarly, experimentation with various interactions effects in the sector attachment equation suggested that interactions between age and age left school produced the best results in terms of the size and significance of the reweighting errors in the RIF decomposition.
3. Matching was performed using 339 unique three-digit occupations from the SOC90 codes (1994–2000), the 353 four-digit occupations from the SOC2000 codes (2001–2010) and the 369 four-digit occupations from the SOC2010 codes (2011–2017).
4. These differences are presented in percentage terms, although technically they are log point differences.

5. Significance tests for these changes are reported in Table 1.
6. The UK recession officially lasted from 2008Q2 to 2009Q2, although unemployment did not peak until 2011Q3 and there were also episodes of negative growth in gross domestic product (GDP) in 2010, 2011 and 2012. In fact, GDP in the UK did not return to its 2008 level until 2014Q3, suggesting that the effects of the recession have been long lasting.
7. Although Bryson and Lucchino (2014) found that mangers in 30 per cent of establishments in the 2011 Workplace Industrial Relations Survey mentioned the NMW as having influenced pay settlements of the largest occupational group.
8. The RIF estimates on which these decompositions are based are available from the authors on request. However, it is worth noting that the estimates share some common features with a standard regression analysis of earnings. Thus, RIF-based estimates tend to increase with age (but at a decreasing rate), educational attainment, occupational status, job tenure and establishment size.
9. The nature of the relationship between the composition and pure composition effects is of interest because the size and significance of the specification error provides less strong support for the linearity assumption made in the RIF-based decomposition.
10. These estimates are available from the authors upon request.
11. On the recommendation of a referee, we also considered removing Higher Education and not-for-profit employees from the analysis. The general pattern of results was the same as specification 2, although the magnitude of the pure composition effect did increase at the 90th percentile for both males and females in the early part of the period. By contrast, differences in the pure structure effect (relative to specification 2) were invariably statistically insignificant.
12. Estimates of these differences and estimated standard errors are available from the authors upon request.

References

Belman, D. and Heywood, J. (1989). 'Establishment size, public administration and government wage differentials'. *Economic Letters*, 29: 95–98.
—— and —— (1990). 'The effect of establishment size and firm size on public wage differentials'. *Public Finance Quarterly*, 18: 221–35.
Bender, K. (1998). 'The central government-private sector wage differential'. *Journal of Economic Surveys*, 12: 177–220.
—— (2003). Examining equality between public and private sector wage distributions. *Economic Inquiry*, 41: 62–79.
—— and Elliott, R. (1999). 'Relative earnings in the UK public sector: the impact of pay reform on pay structure'. In R. Elliott, C. Lucifora and D. Meurs (eds.), *Public Sector Pay Determination in the European Union*. London: Macmillan, pp. 285–339.
Blackaby, D., Felstead, A., Jones, M., Makepeace, G., Murphy, P. and Wass, V. (2015). 'Is the public sector pay advantage explained by differences in work quality?'. In A. Felstead, D. Gallie and F. Green (eds.), *Unequal Britain at Work*. Oxford: Oxford University Press, pp. 147–66.
——, Murphy, P. and O'Leary, N. (1999). 'The payment of public sector workers in the UK: reconciliation with North American findings'. *Economics Letters*, 65: 239–43.
——, ——, —— and Staneva, A. (2018). 'Regional pay? The public/private sector pay differential'. *Regional Studies*, 52: 477–89.

Blinder, A. (1973). 'Wage discrimination: reduced form and structural estimates'. *Journal of Human Resources*, 8: 436–55.

Brown, C. and Medoff, J. (1989). 'The employer size-wage effect'. *Journal of Political Economy*, 97: 1027–59.

Bryson, A., Forth, J. and Stokes, L. (2017). 'How much performance pay is there in the public sector and what are its effects'. *Human Resource Management Journal*, 27: 581–97.

—— and Lucchino, P. (2014). 'The Influence of the National Minimum Wage on Pay Settlements in Britain'. National Institute of Economic and Social Research, 1–61.

Cribb, J. and Emmerson, C. (2014). 'Workplace Pensions and Remuneration in the Public and Private Sectors in the UK'. IFS Briefing Note, No. BN151. London: Institute of Fiscal Studies.

Danzer, A. and Dolton, P. (2012). 'Total reward and pensions in the UK in the public and private sectors'. *Labour Economics*, 19: 584–94.

Depalo, D., Giordano, R. and Papapetrou, E. (2015). 'Public-private sector wage differentials in euro-area countries: evidence from quantile decomposition analysis'. *Empirical Economics*, 49: 985–1015.

DiNardo, J., Fortin, N. and Lemieux, T. (1996). 'Labour market institutions and the distribution of wages, 1973–1992: a semi parametric approach'. *Econometrica*, 64: 1001–44.

Disney, R., Emmerson, C. and Tetlow, G. (2009). 'What is a public sector pension worth?'. *Economic Journal*, 100: 91–104.

—— and Gosling, A. (1998). 'Does it pay to work in the public sector?'. *Fiscal Studies*, 19: 347–74.

Doeringer, P. and Piore, M. (1971). *Internal Labor Markets and Manpower Analysis*. Lexington, MA: Heath Lexington Books.

Dolton, P., Makepeace, G. and Marcenaro-Guitierrez, O. (2015). 'Public sector pay in the UK: quantifying the impact of the review bodies'. *Manchester School*, 83: 701–24.

Elliott, R. and Bender, K. (1997). 'Decentralisation and pay reform in central government: a study of three countries'. *British Journal of Industrial Relations*, 35: 447–75.

Emmerson, C., Johnson, P. and Joyce, R. (2016). 'IFS Green Budget 2016'. Institute for Fiscal Studies, London.

Firpo, S., Fortin, N. and Lemieux, T. (2007). 'Decomposing wage distributions using recentred influence function regressions'. Mimeo, University of British Columbia.

——, —— and —— (2018). 'Decomposing wage distributions using recentred influence function regressions'. *Econometrics*, 6: 1–40.

Fogel, W. and Lewin, D. (1974). 'Wage determination in the public sector'. *Industrial and Labor Relations Review*, 27: 410–31.

Fortin, N., Lemieux, T. and Firpo, S. (2010). 'Decomposition Methods in Economics'. Working Paper 16045, National Bureau of Economic Research.

Gregory, M. (1990). 'Public sector pay'. In M. Gregory and A. Johnson (eds.), *A Portrait of Pay 1970–1982: An Analysis of the New Earnings Survey*. Oxford: Oxford University Press, pp. 172–205.

Gregory, R. and Borland, J. (1999). 'Recent development in public sector labour markets'. In O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, Vol. 3C. New York and Oxford: Elsevier Science, North-Holland, pp. 3573–630.

Jones, F. (1983). 'On decomposing the wage gap: a critical comment on Blinder's method'. *Journal of Human Resources*, 18: 126–30.

Lausev, J. (2014). 'What has 20 years of public/private sector pay gap literature told us? Eastern European transitioning vs. developed economies'. *Journal of Economic Surveys*, 28: 516–50.

Low Pay Commission (2012). 'National minimum wage'. Low Pay Commission Report, Cm 8302.

Lucifora, C. and Meurs, D. (2006). 'The public sector pay gap in France, Great Britain and Italy'. *Review of Income and Wealth*, 52: 43–60.

Machado, J. and Mata, J. (2005). 'Counterfactual decomposition of changes in the wage distribution using quantile regression'. *Journal of Applied Econometrics*, 20: 445–65.

Maczulskij, T. (2013). 'Public-private sector wage differentials and the business cycle'. *Economic Systems*, 37: 284–301.

Melly, B. (2005). 'Public-private sector wage differentials in Germany: evidence from quantile regression'. *Empirical Economics*, 30: 505–20.

Millard, B. and Machin, A. (2007). 'Characteristics of public sector workers'. *Economic and Labour Market Review*, 1: 45–55.

Miller, P. (2009). 'The gender pay gap in the US: does sector make a difference?'. *Journal of Labor Research*, 30: 52–74.

Moulton, B. (1990). 'A re-examination of the federal-private wage differential in the United States'. *Journal of Labor Economics*, 8: 270–93.

Mueller, R. (1998). 'Public-private sector wage differentials in Canada: evidence from quantile regressions'. *Economics Letters*, 60: 229–35.

Nickell, S. and Quintini, G. (2002). 'The consequences of the decline in public sector pay in Britain: a little bit of evidence'. *Economic Journal*, 112: F107–18.

Oaxaca, R. (1973). 'Male-female wage differentials in urban labor markets'. *International Economic Review*, 14: 693–709.

Postel-Vinay, F. and Turon, H. (2007). 'The public pay gap in Britain: small differences that (don't?) matter'. *Economic Journal*, 117: 1460–503.

Poterba, J. and Rueben, K. (1994). 'The Distribution of Public Sector Wage Premia: Evidence Using Quantile Regression Methods'. Working Paper No. 4734, National Bureau of Economic Research.

Propper, C. and Van Reenen, J. (2010). 'Can pay regulation kill? Panel data evidence on the effects of labour markets on hospital performance'. *Journal of Political Economy*, 118: 222–73.

Pyper, D., McGuiness, F. and Brien, P. (2018). 'Public Sector Pay'. Briefing Paper, CBP 8037, House of Commons Library, pp. 1–18.

Rees, H. and Shah, A. (1995). 'Public-private sector wage differentials in the UK'. *Manchester School*, 63: 52–68.

Runge, J., Hudson-Sharp, N. and Rolfe, H. (2017). *Use of Agency Workers in the Public Sector: Report to OME*. London: NIESR.

School Teachers' Review Body (2016). *Twenty-Sixth Report 2016 Cm9302*. London: Office of Manpower Economics.

Stewart, M. (2012). 'Wage inequality, minimum wage effects, and spillovers'. *Oxford Economic Papers*, 64: 616–34.

Topel, R. (1991). 'Specific capital, mobility and wages: wages rise with job seniority'. *Journal of Political Economy*, 99: 145–76.

Yu, K., Van Kerm, P. and Zhang, J. (2005). 'Bayesian quantile regression: an application to the wage distribution in 1990s Britain'. *Indian Journal of Statistics*, 67: 359–77.

Appendix

FIGURE A1
Public-Sector Employment, 1994–1997: LFS and PSE.

TABLE A1
Variable Definitions

	Definition
Real hourly pay	The natural logarithm of gross hourly earnings defined in January 1990 prices.
Age	Age of the respondent in year, entered in linear and quadratic form.
Years of Schooling	The number of years spent in full-time education.
Qualifications	A series of dummy variables denoting the highest equivalent educational qualification of the respondent: 1 — degree; 2 — A-Level; 3 — O-Level/GCSE; 4 — other; 5 — none (E).
Job tenure	A series of dummy variables denoting the tenure with current employer of the respondent: 1 — ≤ 1 year; 2 — 1–2 years; 3 — 2–5 years; 3 — 5–10 years; 4 — 10–20 years; 5 — 20+ years.
Marital status	A series of dummy variables denoting the marital status of the respondent: 1 — married; 2 — single (E); 3 — widowed, divorced or separated.
Region of work	A series of dummy variables denoting the government office region of work of the respondent: 1 — Northern (E); 2 — Yorkshire & Humberside; 3 — East Midlands; 4 — East Anglia; 5 — London; 6 — South East; 7 — South West; 8 — West Midlands; 9 — North West; 10 — Wales; 11 — Scotland.
White	A dummy variable that takes on the value of 1 if the respondent is of a white ethnic background, and 0 otherwise.
Occupation	A series of dummy variables denoting the occupation of the respondent: 1 — higher managerial and professional; 2 — lower managerial and professional; 3 — intermediate occupations; 4 — lower supervisory and technical (E); 5 — semi-routine occupations; 6 — routine occupations.
Establishment size	A series of dummy variables denoting the establishment size where the respondent works: 1 — 1–10 employees; 2 — 11–20 employees; 3 — 21–24 employees; 4 — 25–49 employees (E); 5 — 50 or more employees.
Number of children	The number of dependent children in the respondent's household.

Note: (E) denotes an excluded dummy variable category.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.