



City Research Online

City, University of London Institutional Repository

Citation: Bryson, A., Forth, J. ORCID: 0000-0001-7963-2817 and Stokes, L. (2017). How much performance pay is there in the public sector and what are its effects?. Human Resource Management Journal, 27(4), pp. 581-597. doi: 10.1111/1748-8583.12153

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/20904/>

Link to published version: <http://dx.doi.org/10.1111/1748-8583.12153>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Full Title: How Much Performance Pay is there in the Public Sector and What Are Its Effects?

Running Title: Performance Pay in the Public Sector

March 2017

Authors

Professor Alex Bryson (corresponding author): a.bryson@ucl.ac.uk; UCL Department of Social Science, Institute of Education, 20 Bedford Way, London, WC1H 0AL

Mr John Forth: j.forth@niesr.ac.uk; National Institute of Economic and Social Research, 2 Dean Street, London, SW1P3HE

Ms Lucy Stokes: l.stokes@niesr.ac.uk; National Institute of Economic and Social Research, 2 Dean Street, London, SW1P3HE

Abstract

Theory suggests performance pay (PP) can align employees' interests with those of the employer, attract high-ability workers and incentivise effort, but that it may be less effective in the public sector. However, empirical evidence on its incidence and effects is largely confined to the private sector. We find half the 20 percentage point gap in PP incidence in Britain between the public and private sectors is accounted for by differences in occupational composition. The gap falls to 8 percentage points when "matching" employees in both sectors on their demographic and job characteristics. PP is linked to positive job attitudes among private sector employees, but not among observationally equivalent private sector employees. Furthermore, PP is negatively correlated with workplace performance in the public sector. These findings raise important questions about public policies promoting PP in the public sector.

Key words

public sector; performance pay; workplace performance; organisational commitment; intrinsic job satisfaction.

Funding

This work is funded by the Economic and Social Research Council, Grant Reference ES/I035846/1.

Acknowledgements

We thank seminar participants at Durham Business School for comments. The authors acknowledge the Department for Business, Energy and Industrial Strategy, the Economic and Social Research Council, the Advisory, Conciliation and Arbitration Service and the National Institute of Economic and Social Research as the originators of the 2011 Workplace Employee Relations Survey data, and the Data Archive at the University of Essex as the distributor of the data.

1. Introduction

Performance pay (PP) is well-established in the private sector in Britain and elsewhere (European Foundation, 2007). It is an important component in many employees' compensation, not just those at the top end of the wage distribution (). It is used to align the interests of employers (principals) and employees (agents), particularly where monitoring effort is costly, and is used both to attract highly able employees (Lemieux et al., 2009) and incentivise effort (Lazear, 1986; Prendergast, 1999). However, traditionally PP has been less prevalent in the public sector (Prentice et al., 2007). This is despite the fact that using PP to improve public sector efficiency has been an important component in governments' public service modernisation agenda for some time, irrespective of their political persuasion (Burgess and Ratto, 2003). More broadly, incentivising public servants through PP is viewed as complementary to the target-setting ethos espoused in New Public Management (NPM) in many developed countries (Bach et al., 2009; Barzelay, 2001).

Performance-based contracts have been used to deliver public services through private sector and third sector providers for some years (Rolfe et al., 1996), and public services across the developed world are increasingly judged as "successes" or "failures" according to whether they hit targets such as waiting times for hospitals and student performance in schools. In Britain there are renewed efforts on the part of government to promote PP across broad swathes of the public sector. For instance, in September 2013 pay progression based on performance was rolled out to all teachers in England and Wales for the first time, at the same time that the mandatory pay scale began to be dismantled (STRB, 2013).

Despite policy interest in the use and value of PP in the public sector the issue has attracted little attention in the academic literature. Two recent reviews uncovered few papers on PP in

the public sector in the UK, but shed some light on the issue of PP in the public sector in other countries (Prentice et al., 2007; Ray et al., 2014). There appears to have been little additional research since then.

We review this literature and contribute to it by considering why PP is less common in the public sector, what its implications are for employees' job attitudes and the performance of the public sector, and what the implications of more PP might be for the delivery of public services. In doing so we draw on a variety of theoretical perspectives. First, using principal-agent theory, we argue that jobs in the public sector have attributes that militate against the use of PP. Second, we suggest that worker preferences are heterogeneous across the public and private sectors such that workers in the public sector may be less inclined to sort into PP jobs and, conditional on being in a PP job, will be less responsive to PP than their private sector counterparts. Third, we suggest public sector employers rely on career incentives to elicit worker effort, as opposed to more short-term incentives like PP. Fourth, we argue unions use their relative strength in the public sector to block widespread use of PP in favour of a fixed rate for the job. Fifth, we argue that organisational benefits of PP will be less evident in the public sector than in the private sector because it is unable to "work" via improvements in job attitudes.

We exploit nationally-representative linked employer-employee data for Britain from the 2011 Workplace Employment Relations Survey (WERS) to examine the incidence and correlates of PP in the private and public sectors. Using theoretical insights from the approaches noted above we test hypotheses as to why we might expect less use of PP in the public sector compared to the private sector and consider the implications of PP for employees and employers.

The remainder of the paper is structured as follows. In Section Two we discuss the literature to date and hypothesise why it is that PP appears less common in the public sector, and why we might expect PP to have different effects on employees in the public sector compared to the private sector, and why its effects on workplace performance may be less evident in the public sector. Section Three introduces the data used in the study and the key measures used in the analysis before presenting our estimation strategy. Section Four presents our results. Section Five summarises the findings and discusses their implications, the limitations of the study and future research avenues before concluding with some remarks on the wider implications of the study.

2. Literature and Hypotheses

In this section we review the literature on differences in the relative incidence of PP in the public and private sectors, and the effects of PP on job attitudes and workplace performance, identifying hypotheses that are testable with our data.

2.1 Explaining the incidence of performance pay in the public and private sectors

PP is far more prevalent in the private sector than it is in the public sector. A number of potential explanations have been proffered in the literature. Most adopt a principal-agent framework in which the employer, as the principal, considers the costs and benefits of incentivising employees (the agent) using incentive pay as opposed to a standard fixed pay contract in order to align the interests of principals and agents in a scenario where monitoring employee effort is costly (Lazear, 2000). In this context, whether it is sensible for an employer to pay for performance depends, in part, on the nature of the tasks performed by the worker. In particular, PP is more likely where monitoring effort is

costly and where the occupation entails a single task producing easy-to-measure outputs. These conditions are less common in the public sector because output is often hard to measure due to the complexity of the public sector “good” or the need to achieve multiple objectives (multi-tasking) (Burgess and Ratto, 2003; Prentice et al., 2007; Burgess et al., 2010).

In our data, described in Section Three, we do not directly observe the degree to which individuals undertake multi-tasking jobs, nor the difficulties the employer faces in measuring output. Instead we observe the (three-digit) occupations they perform. We conceive of occupations as bundles of tasks which capture some of the heterogeneity in the nature of work (although there is much within-occupation variance in the nature of work which we are not capturing). Other research (Bryson and Forth, forthcoming), and the analysis reported below, indicates that only a sub-set of occupations exist in both sectors. It is therefore conceivable that differences in the occupations performed in the two sectors, and thus the mix of tasks performed by workers in those sectors, will affect the incidence of PP. Thus, conditioning on the sectoral mix in occupations will reduce the PP gap between the public and private sectors that arises from differences in the nature of work tasks performed in the two sectors. We therefore test the following hypothesis:

Hypothesis One: The PP incidence "gap" between the public and private sectors will narrow when one compares public and private sector employees in the same occupations.

Heterogeneous employee preferences across the two sectors may also affect employees’ desire for PP. For example, because public sector jobs are typically more protected, they tend to attract workers who are relatively more risk averse (Pfeifer, 2011) and, as such, less willing to accept PP (Alesina et al., 2001). If the employer has to pay a risk premium to compensate more risk averse workers for the potential risks of income variability associated with PP, this will raise the costs of PP relative to fixed pay, reducing the viability of PP from the principal’s perspective. We do not observe risk preferences

in our data. However, these (to us) unobserved characteristics may be correlated with those traits we do observe such that we capture some of this selection, at least in part. Thus, for instance, there is a large literature on gender and risk preferences, such that matching on gender partially captures this aspect of worker selection into the public sector (Dohmen and Falk, 2011). This leads us to our second hypothesis:

Hypothesis Two: The PP incidence "gap" between the public and private sectors will narrow when one compares observationally equivalent employees in the two sectors

A third potential reason for the PP “gap” between public and private sector employees relates to the strength of trade unions in the public sector. In Britain, unionisation rates are much higher in the public sector than they are in the private sector, whether measured in terms of union recognition for pay bargaining or the percentage of employees covered or in membership (van Wanrooy et al., 2013). In bargaining for better terms and conditions for union-covered workers, unions have traditionally argued for a "rate-for-the-job", that is, a standardised wage for a job, thus limiting employers' ability to permit wages to vary with worker performance (Gittleman and Pierce, 2013; O'Halloran, 2013). To the extent that union strength helps unions achieve their bargaining objectives, we expect them to limit the use of PP in the public sector relative to the private sector, leading to our third hypothesis:

Hypothesis Three: The PP incidence gap between the public and private sectors will narrow when differences in unionisation across the sectors are accounted for.

2.2 Employee and organisational outcomes of PP

As noted above, employers use PP to incentivise worker effort (by rewarding workers for at least some of their additional output) and to attract more able workers by offering them the opportunity to earn a wage that better reflects their actual productivity. Using PP to attract more able employees may be profitable for the employer even in the absence of incentive effects. In both cases, however, the underlying assumption is that employees respond well to financial incentives. This appears to be the case among private sector employees. For instance, PP is associated with greater job satisfaction in the private sector (Bryson et al., 2016).¹ PP also engenders greater organisational commitment (OC) among private sector employees, either as part of a reciprocal exchange between employer and employees (Bryson and Freeman, 2014), or simply because PP permits the more able employees who seek it out to earn more from the output their efforts produce. However, public sector employers have traditionally relied upon a good total reward package, including pensions, to attract high calibre candidates (Danzer and Dolton, 2012) and have used career incentives based on promotion opportunities, rather than PP, to incentivise workers (Prendergast, 1999).

We can conceive of workers as heterogeneous in their tastes for particular jobs and working environments. Besley and Ghatak (2005) draw a distinction between profit-oriented and mission-oriented production. The latter, which entails the production of collective goods, characterises much of the public sector and attracts worker types who wish to pursue goals because they perceive intrinsic benefits from doing so, as opposed to simply maximising income from their activities. The implication is that those attracted by mission-oriented work, such as that in the public sector, attach greater value to intrinsic job rewards, such as a sense of achievement, than they might do to extrinsic rewards such as pay. This makes them particularly responsive to ‘intrinsic’ rewards such that PP may not be required for them to perform well (Besley and Ghatak, 2005). Indeed, pecuniary incentives may even be counter-productive where they reduce public servants' intrinsic job satisfaction (Burgess and Metcalfe, 2000). We therefore test the following two related hypotheses:

¹ Green and Heywood (2008) also find a positive association in worker fixed effects models which pool private and public sector workers.

Hypothesis Four (a): PP is positively associated with intrinsic job satisfaction in the private sector, whereas the association is non-significant or negative in the public sector.

Hypothesis Four (b): PP is positively associated with organisational commitment in the private sector, whereas the association is non-significant or negative in the public sector.

Those advocating more extensive use of PP in the public sector start from the premise that doing so will be beneficial in terms of improvements in public service provision, either because it will assist in attracting high quality workers to the public sector, or because it will incentivise worker effort. Either way, if PP is currently below its optimal level in the public sector, one would expect public sector workplaces with PP to have higher productivity levels than "like" public sector workplaces without PP. In their review of the literature Prentice et al (2007) find tentative evidence that PP schemes do bring benefits in the public sector, but some of the effects have been limited by the design of schemes and "gaming", while many studies are based on non-representative studies. However, since then some studies have emerged to suggest that PP can positively affect workplace and organisational performance in the public sector. Burgess et al. (2010) show the introduction of team-based PP improved task allocation in HM Customs and Excise, resulting in improved team performance. Imberman (2015), in his review of PP effectiveness for teacher performance around the world, finds PP can improve teacher performance, but that effects depend on the design of the PP scheme. We revisit this issue by looking at performance outcomes for the representative workplace data described in Section Three. To our knowledge ours is the only study that seeks to establish the relationship between PP and performance in the public sector using nationally representative data. In the light of the existing limited evidence regarding PP and public sector performance the fifth and final hypothesis we test is:

Hypothesis Five: Public sector workplaces with PP will have higher reported performance than similar public sector workplaces without PP.

3. Methods

In this section we introduce our data, present the key measures used in our analyses, and describe our estimation strategy.

3.1 Data

We analyse linked employer-employee data from the Workplace Employment Relations Survey 2011 (WERS). Appropriately weighted, it is a nationally representative survey of workplaces in Britain with 5 or more employees covering all sectors of the economy except agriculture and mining (van Wanrooy et al., 2013). The analysis exploits two elements of the survey. The first is the management interview, conducted face-to-face with the most senior workplace manager responsible for employee relations. Interviews were conducted in 2,680 workplaces between March 2011 and June 2012 with a response rate of 46%. The second element is the survey of employees, distributed in workplaces where a management interview was obtained. Self-completion questionnaires were distributed to a simple random sample of 25 employees (or all employees in workplaces with 5-24 employees) in the 2,170 workplaces where management permitted it. Of the 40,513 questionnaires distributed, 21,981 (54%) usable ones were returned.² Weights are provided with the survey data to correct for the sample design and any observable non-response biases.

² An additional 3,858 questionnaires were distributed at 247 workplaces where there were no employee questionnaires returned. It is assumed that these questionnaires were never distributed by the employer (van Wanrooy et al., 2013: 210) so they are not included in the figures in the text.

To identify the occupational overlap in the public and private sectors we use the 2011 and 2012 Quarterly Labour Force Survey. For each Unit Group of the 2010 Standard Occupational Classification we identify the proportion of all employees in each four-digit occupation who report that they are employed in the public sector, excluding any occupations for which this proportion is outside a specified range.

3.2 Measures

Performance pay: these measures are taken from the employee and employer questionnaires. Employees are asked "Which of the following do you receive in your job here...Payments based on your individual performance or output; payments based on the overall performance of a group or a team; payments based on the overall performance of your workplace or organisation (eg. profit-sharing scheme)". They are instructed to tick all that apply.³ We can therefore distinguish between employees' receipt of PP arising from individual, team and organisational performance. These types of performance pay have different implications for the way employees are paid, such that studies which are unable to make these distinctions may be conflating types of performance pay scheme, or else capturing just one of them, making it difficult to interpret results or to compare results across studies.⁴ In establishing the link between PP and workplace performance we characterise public sector workplaces according to the employer's definition of whether there is any PP use at the workplace. We derive a dummy variable from a series of questions managers are asked regarding the incidence of various PP schemes at the workplace: "Do any employees in this workplace get paid by results or receive merit pay; profit-related payments or profit-related bonuses; operate any of the employee share schemes listed on this card for any of the employees at this workplace?"⁵

³ The question also includes the following aspects of compensation: basic fixed/salary wage; extra payments for additional hours of work or overtime; contributions to a pension scheme.

⁴ For a discussion of various types of performance pay schemes and analyses pointing to differences in the correlates of different types of scheme see Bryson et al. (2013).

⁵ Employees of Royal Mail account for the vast majority of public sector employees covered by a share scheme <http://www.royalmailgroup.com/investors/shareholder-communications/managing-your-shares>

Work attitudes: We follow White and Bryson (2013: 393-395) in the construction of two dependent variables. The scale for intrinsic job satisfaction (IJS) combines 5-point Likert scales from "very satisfied" (5 points) to "very dissatisfied" (1 point) for four job facets, namely "sense of achievement you get from your work", "the scope for using your own initiative", "the amount of influence you have over your job", and "the work itself". The scale for organisational commitment (OC) combines 5-point Likert scales from "strongly agree" (5 points) to "strongly disagree" (1 point) on three items, namely "I share many of the values of my organisation", "I feel loyal to my organisation" and "I am proud to tell people who I work for". They have scale reliability coefficients (Cronbach Alphas) of 0.87 and 0.85 respectively. Both scales are converted into z-scores based on the whole sample (public and private sector employees) with a mean of zero and standard deviation of one so we can compare the quantitative effect of PP across the two outcomes.

Workplace performance: this is measured using the manager's subjective assessment on three separate measures.⁶ We follow Bryson et al. (forthcoming) in the construction of the dependent variable. It is an additive scale combining managers' responses to three questions: "Compared to other workplaces in the same industry how would you assess your workplace's...financial performance; labour productivity; quality of product or service". Responses are recorded on a 5-point Likert scale from "a lot better than average" to "a lot below average". The "a lot below average" and "below average" codes are collapsed and scales scored from 0 to 3 where 3="a lot above average". Summing them gives a scale of 0 ('below average' performance on all three items) to 9 (performance 'a lot better than average' on all 3 items). Among public sector workplaces (the focus for our subsequent analysis), the pairwise correlations between the three measures vary between 0.35 (financial performance

⁶ These measures are frequently used in the literature. For a recent example see Wu et al. (2015). For a discussion of these measures and their relationship with accounting measures of performance see Forth and McNabb (2008).

and product/service quality) and 0.52 (financial performance and labour productivity). Factor analysis identifies a single factor with an eigen value of 1.25, and the alpha reliability coefficient for the composite performance scale is 0.72.

Public sector: Workers are often not well informed about the public sector status of their employer (see Office for National Statistics, 2010: 6), leading to measurement error when distinguishing public and private sector employees. WERS relies on the managerial respondent to identify the status of the workplace, thus minimising this measurement error problem (Blanchflower and Bryson, 2010).

3.3 Estimation Strategy

To establish what role occupational differences make to PP incidence in the public and private sectors we focus on employees in similar or identical occupations using different matching methods. We then assess the role of other worker traits on the PP “gap” in the two sectors using propensity score matching (PSM), described in more detail below, to compare PP incidence in the two sectors among employees who are observationally equivalent. We then use PSM again to compare PP employees with observationally equivalent fixed pay employees within the public and private sectors to estimate the association between PP and employee job attitudes in each sector. Finally we use PSM to compare PP with observationally equivalent non-PP workplaces within the public sector to establish the effect of PP on workplace performance in the public sector.

To see how much of the PP “gap” is accounted for by occupational differences across sectors (Hypothesis One) we run employee-level analyses on the sub-set of employees in occupations that appear in both the public and private sectors in the Labour Force Survey. We thus focus on jobs that are undertaken in both sectors. We then select all of the employees in the WERS

data set who work in these 'cross-sectoral' occupations and use these to estimate the incidence of PP in the public and private sectors. We vary the exclusion criterion in order to check the sensitivity of the PP incidence estimates to our choice of upper and lower bounds.

Our second approach is to "hard match" employees across sectors on their four-digit occupation. This entails identifying the occupations undertaken by the public sector employees in WERS and matching them to private sector employees from the same occupation. This ensures that we obtain sub-samples of private and public sector employees with identical occupational profiles.

Our third approach is to match employees according to their propensity to be a public sector employee, conditional on their observable characteristics. This is known as propensity score matching (PSM). We run a probit estimator on all employees in WERS to generate an estimated propensity for being a public sector employee that lies between zero and one. The observable traits we use to match on are gender, age, education, region, industry and occupation. The resulting propensity score is used to match public and private sector employees, reweighting private sector employees such that they approximate the public sector employees in the distribution of their observable characteristics, not just their occupation. Where the estimated propensity score for public sector employees is too distant from any of the scores in the private sector those public sector employees are omitted from the matched sample. The assumption is that, having reweighted the private sector employees so that their observable traits appear similar to those in the public sector, the outcome - in this case receiving PP - that public sector employees would have received had they been in the private sector is independent of their status as a public sector employee. In this way we construct a third sample of matched employees, and identify the relative incidence of PP in

the two sectors having accounted for a range of key differences between the workers who enter the two sectors, and attribute the residual difference in the prevalence of PP to being in the public sector. Since the matching estimator conditions on worker demographic traits as well as occupation this allows us to test Hypothesis Two, namely that some of the PP differential across sectors is due to the characteristics of employees in the two sectors. We return to the details of how we perform the PSM below.

To establish the extent to which unionisation accounts for lower PP incidence in the public sector (Hypothesis Three) we compare PP incidence among employees in the PSM-matched sample, but this time we add employees' union coverage into the probit generating the propensity scores. In this way we ensure that the public/private sector comparison accounts for differences in union coverage across the sectors, thus approximating a counterfactual world in which unionisation is no different across the two sectors.

We assess the PP effect on two job attitudes (Hypothesis Four) - organisational commitment (OC) and intrinsic job satisfaction (IJS) - in the two sectors by comparing differentials in the IJS and OC of employees who receive PP and their matched equivalents who do not receive PP. We do this separately for the public and private sectors.

Hypothesis Five is that public sector workplaces will perform better in the presence of PP. First we compare workplace performance among public sector workplaces with and without PP. Then we rerun the comparison using a matching estimator which re-weights the non-PP public sector workplaces such that they appear similar on their observable characteristics to the public sector workplaces using PP. The covariates in the matching estimator were sector, region, workplace size, whether a single-workplace organisation, age of workplace,

recognition of a trade union for pay bargaining, and largest non-managerial occupational group at the workplace. We then recover the mean differences in workplace performance from this matching exercise. This provides an estimate of the PP effect on public sector workplace performance for those who have PP (known in the literature as the average treatment effect on the treated, ATT). However, an additional interesting parameter for policy purposes is the effect that PP would have on the performance of those public sector workplaces that are currently without PP. This is the average treatment effect on the non-treated (ATNT): it captures the effect PP would have on workplace performance among the sample of public sector workplaces that do not use it currently. Recovering this effect is relatively straightforward: it entails running the probit estimator for the probability of PP for public sector workplaces, then using these probabilities to identify from among the PP workplaces the matched comparators to the public sector workplaces without PP.

4. Results

[INSERT TABLE ONE]

Almost one-quarter (23 per cent) of employees say they are paid for performance but the incidence of PP is four-times higher in the private sector than the public sector (27 per cent compared with 7 per cent) and it is higher whether performance is assessed at individual, group/team or workplace/organisation level.⁷ The raw gap in PP incidence between employees in the public and

⁷ This is the case even if one ignores financial participation through share plans and profit sharing. Opportunities for financial participation are necessarily limited in the public sector. Only a small part of the public sector produces tradeable goods or services from which a profit might be generated. Instead, most public sector employers deliver services which are heavily subsidised by tax-payers and are free at the point of use. Public service providers are owned by government so shares in public sector employers are rarely traded or made available to employees. By contrast, private sector employers are free to share equity with their employees in

private sectors is 20 percentage points (see the first pair of columns in Table 1, comparison 1). However, this falls to 14 percentage points when we confine the comparison to employees in those occupations where at least 25 per cent of all employees nationally are in the public sector and at least 25 per cent are in the private sector (comparison 3). The gap falls still further to 10 percentage points (comparison 5) - half the raw gap - when we "hard match" on occupation by focusing attention only on public sector employees in occupations for whom there is a private sector occupational match in WERS. In fact, the public-private differential in PP roughly halves whether performance is measured at the individual (10 to 5 percentage points), team (7 to 3 percentage points) or workplace/organisation level (13 to 7 percentage points). Occupational matching therefore accounts for a sizeable part of the PP differential across the two sectors, a finding which is consistent with Hypothesis One. That is, the gap in incidence of PP between the private and public sectors narrows when comparing employees in the same occupation.

In the final pair of columns in Table 1 (comparison 6) we present the PP differentials having used propensity score matching to obtain counterfactual employees in the private sector for those in the public sector. The matching was undertaken using a kernel estimator with bandwidth 0.6. 920 respondents in 5 occupations were dropped because there was no variance in public sector status within their 3-digit occupation. Three additional public sector employees were too distant from the private sector employees to find a match so they have been dropped from the analysis. The final matching estimates are therefore based on 7,587 public sector employees and 11,779 private sector employees. The advantage of this approach is that we match on a wider set of covariates than simply occupation. Observable features of public sector employees other than their occupation account for some of the PP gap between them and private sector employees, so that the gap falls still further, consistent with Hypothesis Two. Nevertheless, a sizeable difference in PP receipt remains even among employees matched using propensity scores. Specifically, the incidence of PP among private

the form of share schemes, or to engage in profit-sharing. The majority of large private sector firms in Britain do use share ownership or profit-sharing schemes, and government has historically offered tax breaks to encourage such behaviour (Pendleton et al., 2009).

sector employees is still twice as high as among public sector employees (16 per cent versus 8 per cent), giving an absolute differential of 8 percentage points (compared with a raw differential of 20 percentage points). Differences in the incidence of all forms of PP among employees in the public and private sectors are statistically significant at a 1 per cent confidence level.

Our third hypothesis is that higher levels of union coverage in the public sector act as a break on PP such that, if we compare employees in the two sectors with similar levels of union coverage, this will account for some of the PP gap. To establish whether this is the case we incorporate employees' union status in the propensity score estimate. This makes no difference to the PP gap between public sector employees and their matched comparators: the PP gap is identical to that reported in the final column of Table 1. This is the case whether we use an employee's self-reported union membership status or the workplace manager's report of whether that employee's occupation at the workplace has pay determined by collective bargaining with a recognised trade union.

Further investigation reveals that the bivariate association between PP and unionisation is negative in the private sector: entered into a linear regression estimating the probability of PP receipt, union coverage has a coefficient of -0.068 (t-statistic 2.13). With the introduction of controls⁸, the coefficient falls to -0.029 and becomes statistically non-significant (t-stat=1.13). If one replaces union coverage with union membership, the raw correlation is negative (coefficient of -0.056, t-stat=2.28), and becomes positive but non-significant with controls (coefficient 0.009, t-stat 0.43). In the public sector coverage is positively correlated with PP receipt (coefficient 0.032, t-stat=2.43) but this becomes non-significant once controls are introduced (coefficient 0.011, t-stat=0.81). Although there is a positive and statistically significant association between union membership and PP in the public sector with individual control variables (coefficient 0.025, t-stat=2.72) this becomes non-significant in the presence of workplace fixed effects (coefficient 0.000, t-stat=0.05) indicating that the correlation

⁸ The controls are: female; gender missing; age (7 dummy variables); ethnicity; qualifications (8 dummies); 3-digit occupation.

is driven by fixed unobservable differences between public sector workplaces with and without union members. There is therefore no empirical support for Hypothesis Three.

Our fourth hypothesis is that PP will be positively associated with IJS and OC in the private sector, whereas the association is likely to be non-significant or negative in the public sector. In the private sector, a simple comparison indicates that employees receiving PP have higher IJS and higher OC than those on fixed pay contracts (Table 2, data reported in rows 1 and 3). The IJS differential falls by around one-quarter (compare data reported in rows 1 and 2 in column 3) and the OC differential falls by around one-third (compare data reported in rows 3 and 4 in column 3) having matched PP employees with observationally similar fixed pay employees using propensity score matching, but the differentials remain sizeable and statistically significant.⁹ The situation in the public sector is very different. As shown in data rows 5 and 7 of column 3, in the unmatched samples IJS and OC are significantly lower for those on PP compared to fixed pay employees. But the differential falls markedly and becomes statistically non-significant when comparing PP employees with their matched counterparts (data rows 6 and 8 in column 3). These results are consistent with Hypothesis Four: the IJS and OC of private sector employees rises with PP, but in the public sector it has no such effect.

[INSERT TABLE TWO]

Hypothesis Five maintains that public sector workplaces with PP will perform better than those without PP. Using propensity score matching to identify the effect of PP on performance, we recover

⁹ Employees in receipt of any performance pay are matched to non-PP counterparts in the same sector - private or public sector - using a kernel propensity score matching estimator with a bandwidth of 0.6. In the private (public) sector 40 (260) respondents in 4 (27) occupations were dropped because there was no variance in PP status within their 3-digit occupation. Six (three) additional private (public) sector employees were too distant from the public (private) sector employees to find a match so they have been dropped from the analysis so the final matching estimates are based on 2,772 (561) PP employees and 8,695 (6,420) fixed pay employees in the private (public) sector.

two different treatment parameters.¹⁰ The first is the effect of PP on performance for those workplaces with PP (average treatment-on-the-treated, ATT). The second is the effect that PP would have on public sector workplaces that, currently, do not use PP (average treatment-on-the-non-treated, ATNT). Results are presented in Table 3.

[INSERT TABLE THREE]

A comparison of the mean performance of public sector PP and non-PP workplaces indicates that the performance of PP workplaces is poorer than that of non-PP workplaces, a difference that is statistically significant at a 94 per cent confidence level (row 1, Table 3). Matching reveals that the performance of the non-PP workplaces is depressed by their observable characteristics such that, accounting for these observable differences, the performance differential widens to -0.331 (row 2). Furthermore, when we estimate the effect that PP would have on those public sector workplaces where it is currently absent, we find the negative effect on workplace performance is larger at -0.459.

5. Discussion and Conclusions

Despite the current policy interest in the use and value of PP in the public sector, there is relatively little large-scale, quantitative evidence on its prevalence or correlates among public sector workers. There is also little evidence regarding the effects of PP in the public sector. We contribute to the literature using a nationally representative survey of workplaces and their employees to examine the reasons for the lower incidence of PP in the public sector relative to the private sector in Britain, and

¹⁰ Public sector workplaces using any form of PP are matched to those not using PP using a kernel propensity score matching estimator with a bandwidth of 0.6. Of the 666 public sector workplaces with non-missing data 206 used PP, 460 did not. Four PP workplaces and 10 non-PP workplaces had propensity scores that were too distant from possible comparators so were dropped from the analyses leaving a matched sample of 652 workplaces.

to investigate the association between PP and employee attitudes to their jobs, on the one hand, and the association between PP and workplace performance on the other. In doing so we use matching techniques which, to our knowledge, have not been used in this literature before. We have seven new findings.

First, for the first time using nationally representative workplace data for Britain, we confirm that there is a substantial “gap” in the incidence of PP between employees in the private and public sectors. Of course, the nature of public services means public sector employees and workplaces are less likely to have financial participation (share ownership and profit-sharing). But the lower incidence of PP in the British public sector is apparent across *all* types of PP.

Second, the nature of public sector jobs accounts for a sizeable part of the “gap” in PP between the two sectors, reducing it by about half from 20 to 10 percentage points, confirming hypothesis one. This halving of the gap is apparent all types of PP (individual, team and workplace/organisation based).

Third, accounting for differences in the composition of the workforce in the two sectors accounts for a small additional part of the “gap”, confirming hypothesis two. When matching employees in similar occupations and with similar observable traits, those in the public sector were about 8 percentage points less likely to be in receipt of PP than their private sector counterparts.

Fourth, union coverage is not a break on the use of PP in either the public or the private sector. Thus, contrary to hypothesis three, stronger unionisation in the public sector plays no role in the PP “gap” between the two sectors.

Fifth, a substantial PP “gap” remains having accounted for the occupational and worker composition of the two sectors.

Sixth, in accordance with hypotheses 4a and 4b, PP has a positive effect on intrinsic job satisfaction and organisational commitment among private sector employees but it has no significant effect on public sector employees.

Seventh, contrary to hypothesis five, public sector workplaces using PP perform more poorly than their matched counterparts who do not use PP. Furthermore, our estimates of the average treatment-on-the-non-treated indicate that PP would adversely affect workplace performance in public sector workplaces that do not use it currently.

The study has implications for the analysis and understanding of human resource management and employment relations in the public sector. Using matching techniques that are not commonly used in this literature we are able to demonstrate that heterogeneity in the jobs and workers across the public and private sectors can account for something like half the PP “gap” between the two sectors. Most of this is driven by the different occupations that exist in the two sectors, a finding that is consistent with principal-agent theories regarding the difficulties of implementing PP in scenarios where monitoring output is costly (Lemieux et al., 2009). The remaining “gap” might be viewed as the “pure” public sector effect, one that is to be expected based on existing theory which points to special features of public sector agencies, such as difficulties measuring outputs and the multiplicity of principals (taxpayers, government, etc.) with conflicting interests which limit the value of PP (Dixit, 2002). Dixit (2002: 697) points out that “sometimes these special characteristics explain why these agencies are in the public sector in the first place”.

It is often asserted that trade unions limit the incidence of some forms of PP, particularly in the public sector where they are strongest, because PP militates against union goals such as paying a rate-for-the-job (Freeman, 1982; Gregg and Machin, 1989; Burgess and Metcalfe, 2000). We find no evidence to support this contention. Although there is clear evidence of union effects on wage levels in the public sector (Blanchflower and Bryson, 2010) this does not extend to the nature and incidence of PP. To the extent that PP is less evident in the public sector, this is not due to union constraints on managerial power to implement PP.

The absence of PP effects on workplace performance in the public sector contrasts markedly with a wealth of evidence indicating that it positively affects performance in private sector organisations (eg. Lazear, 2000). One reason why human resource management practices such as PP can improve workplace performance is that they can engender IJS and OC (White and Bryson, 2013, review the literature and evidence on the links between HRM and IJS and OC). More generally, Bryson et al. (forthcoming) find improvements in worker wellbeing are associated with improvements in private sector workplace performance. However, we only find a positive association between PP and IJS and OC among private sector employees, and no association (or a weakly negative effect) among public sector employees, suggesting this channel by which PP may affect workplace performance could operate, but only in the private sector. These findings are consistent with the literature (Besley and Ghatak, 2005; Prentice et al., 2007) that suggests public sector workers may differ in their motivations from workers in the profit-oriented sector.

PP may affect workplace performance, even in the absence of associations with IJS and OC, where it incentivises effort, or where it leads to improvements in the allocation of resources. Recent evaluations have shown that PP can increase output in the public sector through both these mechanisms (Burgess et al., 2010) but our evidence suggests that any such effects are not common or,

if they are, the benefits are offset by the costs, such that no general effect on workplace performance is discernible.

There are limitations to the study which could be addressed in future work. First, we lack data on aspects of employees' jobs that might have given us a better indication of whether some of the factors militating against PP use are, in fact, at work. We do not have good indicators for monitoring costs, nor the degree to which employees are multi-tasking. Instead we use detailed occupational classifications to characterise differences in jobs. These may capture some of the job features which are salient in the principal-agent literature (eg. Holmstrom and Milgrom, 1991) regarding the use of PP, but there will be within-occupation variance in these features which our analysis is unable to detect. Nor do we directly observe workers' preferences, such as those relating to risk-taking, which the literature identifies as important in terms of worker sorting across the public and private sectors, and across PP and fixed pay jobs. Instead, we must rely on matching employees according to the features we observe in the data. We do not know how strongly these features are correlated with the factors deemed salient in the literature, so it may be that we have under-estimated the degree to which worker heterogeneity affects the PP "gap" between the public and private sectors. Unobserved worker heterogeneity may also affect our estimates of the relationship between PP and employees' IJS and OC in the public and private sectors. If we think some of these unobserved worker traits are fixed over time it would be possible to net them out of analyses using panel data on workers using individual fixed effects estimates.

A further limitation of the study is that we cannot definitively state that the relationship between PP and public sector workplace performance is causal. This is because, although we are able to match workplaces on observable characteristics, we cannot discount the possibility that our estimates are biased by unobserved differences across workplaces that affect both PP and performance. Future

research may endeavour to overcome this where there is a natural experiment leading to an exogenous change in the incidence of PP, for instance.

Notwithstanding the study's limitations, it provides robust evidence suggesting there are features of public sector employment that mean it is difficult to motivate employees and improve workplace performance through financial incentives of this type. Our estimates of the likely impact of PP on performance in those public sector workplaces where it is currently absent suggest that the benefits of policies promoting the extension of PP in the public sector are likely to be limited or non-existent. This raises important questions about the future of public sector service provision in Britain and, potentially, elsewhere, given policy interest in encouraging the use of PP, especially among professional employees such as teachers where employers have traditionally relied on career-based incentives.

References

- Alesina, A., Danninger, S. and Rostagno, M. (2001) "Redistribution through public employment: the case of Italy". *IMF Staff Papers*, Vol. 48(3), pp. 447-473
- Bach, S. and Givan, R. K. and Bryson, A. (2009) "The Public Sector in Transition", Chapter 13 in Brown, W., Bryson, A., Forth, J. and Whitfield, K. (eds.) *The Evolution of the Modern British Workplace*, Cambridge University Press
- Barth, E., Bratsberg, B., Haegeland, T. and Raaum, O. (2012) "Performance Pay, Union Bargaining and Within-firm Wage Inequality", *Oxford Bulletin of Economics and Statistics*, 74, 3: 327-362
- Barzelay, M. (2001) *The New Public Management: Improving Research and Policy Dialogue*, Berkeley: University of California Press
- Bell, B. and Van Reenen, J. (2014) "Bankers and their Bonuses", *The Economic Journal*, 124, 574: F1-F21
- Besley, T. and Ghatak, M. (2005) "Competition and Incentives with Motivated Agents", *American Economic Review*, 95(3): 616-636
- Blanchflower, D. G. and Bryson, A. (2010) 'The Wage Impact of Trade Unions in the UK Public and Private Sectors', *Economica*, Vol.77, Issue 305, pp.92-109
- Böckerman, P., Bryson, A. and Ilmakunnas, P. (2013) 'Does High Involvement Management Lead to Higher Pay?', *Journal of the Royal Statistical Society Series A*, 176, 4: 861-885
- Bryan, M. and Bryson, A. (2016) "Has Performance Pay Increased Wage Inequality in Britain?", *Labour Economics*, 41: 149-161
- Bryson, A., Clark, A. E., Freeman, R.B., and Green, C. P. (2016) "Share Capitalism and Worker Wellbeing", *Labour Economics*, 42: 151-158

- Bryson, A., Dorsett, R. and Purdon, S. (2002) *The use of propensity score matching in the evaluation of active labour market policies*, Department for Work and Pensions Working Paper No.4
- Bryson, A. and Forth, J. (Forthcoming) "Wage Growth in Pay Review Body Occupations", Office for Manpower Economics
- Bryson, A., Forth, J. and Stokes, L. (Forthcoming) "Does Worker Wellbeing Affect Workplace Performance?", *Human Relations*
- Bryson, A. and Freeman, R. (2010) 'How does shared capitalism affect economic performance in the UK?', Chapter 6 in D. Kruse, R. Freeman and J. Blasi (eds.) *Shared Capitalism at Work: Employee Ownership, Profit and Gain Sharing, and Broad-based Stock Options*, pp. 201-224, University of Chicago Press
- Bryson, A. and Freeman, R. B. (2014) "Employee Stock Purchase Plans: Gift or Incentive?", *IZA Discussion Paper No. 8537*
- Bryson, A., Freeman, R., Lucifora, C., Pellizzari, M and Perotin, V. (2013) "Paying for Performance: Incentive Pay Schemes and Employees' Financial Participation", in T. Boeri, C. Lucifora and K. J. Murphy (eds.) *Executive Remuneration and Employee Performance-related Pay: A Transatlantic Perspective*, Oxford University Press
- Burgess, S. and Metcalfe, P. (2000) "The Use of Incentive Schemes in the Public and Private Sectors: Evidence from British Establishments", *CMPO Working Paper Series No. 00/15*
- Burgess, S., Propper, C., Ratto, M., von Hinke Kessler Scholder, S. and Tominey, E. (2010) "Smarter Task Assignment or Greater Effort: The Impact of Incentives on Team Performance", *The Economic Journal*, 120, 547: 968-989
- Burgess, S. and Ratto, M. (2003) "The Role of Incentives in the Public Sector: Issues and Evidence", *Oxford Review of Economic Policy*, 19, 2: 285-300

- Coleman, T. (2010) "Do financial incentives for delivering health promotion counselling work? Analysis of smoking cessation activities stimulated by the quality and outcomes framework" *BMC Public Health*, 10, 167-172
- Cornelissen, T., Heywood, J. and Jirhahn, U. (2010) "Performance Pay, Risk Attitudes and Job Satisfaction", *Labour Economics*, 18, 2: 229-239
- Danzer, Alexander M and Dolton, Peter J (2012) "Total reward and pensions in the UK in the public and private sectors", *Labour Economics*, 19 (4). pp. 584-594
- Dixit, A. (2002) "Incentives and Organizations in the Public Sector: An Interpretative Review", *Journal of Human Resources*, 37:4, 696-727
- Dohmen, T. and Falk, A. (2011) "Performance pay and multidimensional sorting: Productivity, Preferences and Gender", *American Economic Review*, 101(2): 556-90
- European Foundation for the Improvement of Living and Working Conditions,(2007) *Fourth European Working Conditions Survey*, Eurofound, Dublin
- Forth, J., Bryson, A. and Stokes, L. (2016) "Are Firms Paying More For Performance?", *International Journal of Manpower*, 37(2): 323-343.
- Forth, J., and McNabb, R. (2008) "Workplace performance: a comparison of subjective and objective measures in the 2004 Workplace Employment Relations Survey", *Industrial Relations Journal*, 39(2), 104–123
- Freeman, R. (1982), "Union Wage Policies and Wage Dispersion within Establishments", *Industrial and Labor Relations Review*, 36:1, pp.3-21
- Gittleman, M. and Pierce, B. (2013) "How Prevalent is Performance-Related Pay in the United States? Current Incidence and Recent Trends", *National Institute Economic Review*, 226: R4-R16
- Green, C.P and Heywood, J.S. (2008) "Does performance pay increase job satisfaction?", *Economica*, 75(300), 710-728

- Gregg, P. A. and Machin, S. J. (1989) "Unions and the incidence of performance linked pay schemes in Britain", *International Journal of Industrial Organization*, vol. 6(1), pages 91-107
- Holmstrom, B. and Milgrom, P. (1991) "Multi-task Principal-Agent Analyses: Incentive Contracts, Asset Ownership and Job Design", *Journal of Law, Economics and Organization*, 7: 24-52
- Imberman, S. A. (2015) "How effective are financial incentives for teachers?", *IZA World of Labor*, 158
- Lazear, E. (1986), "Salaries and piece rates", *Journal of Business*, 59: 3, 405-31
- Lazear, E. (2000) "Performance Pay and Productivity", *American Economic Review*, 90, 1346-1361
- Lemieux, T., MacLeod, W. B., Parent, D. (2009) "Performance Pay and Wage Inequality" *Quarterly Journal of Economics*, 124 (1): 1–49
- O'Halloran, P. (2013) "Union coverage, membership and performance-related pay: are piece rates different?", *National Institute Economic Review* No. 226: R30-R41
- Office for National Statistics (2009) *Summary Quality Report for the Annual Survey of Hours and Earnings*, Newport: ONS
- Office for National Statistics (2010) *Public Sector Employment: Quality and Methodology Information*, Newport: ONS.
- Pendleton, A., Whitfield, K. and Bryson, A. (2009) 'The changing use of contingent pay in the modern British workplace', Chapter 11 in W. Brown, A. Bryson, A., Forth, J. and K. Whitfield (eds.) *The Evolution of the Modern Workplace*, pp. 256-284, Cambridge University Press
- Pfeifer, C. (2011) "Risk Aversion and Sorting into Public Sector Employment", *German Economic Review*, 12: 85–99
- Prendergast, C. (1999) 'The provision of incentives in firms' *Journal of Economic Literature*, 37, 1: 7

- Prentice, G., Burgess, S. and Propper, C. (2007) *Performance pay in the public sector: A review of the issues and evidence*, Office of Manpower Economics, London
- Ray, K., Foley, B., Tiffany, T., Walne, D. and Bajorek, Z. (2014) *A review of the evidence on the impact, effectiveness and value for money of performance-related pay in the public sector*, The Work Foundation
- Rolfe, H., Bryson, A. and Metcalf, H. (1996) *The Effectiveness of TECs in Achieving Jobs and Qualifications for Disadvantaged Groups*, Department for Education and Employment Research Studies RS4, HMSO, London
- Rosen, S. (1990) 'Contracts and the market for executives', *NBER Working Paper No. 3542*
- School Teachers Review Body (2013): Twenty-Second Report – 2013 (download: https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/208618/School_Teachers_Review_Body_-_22nd_report.pdf)
- van Wanrooy, B., Bewley, H., Bryson, A., Forth, J., Freeth, S., Stokes, L. and Wood, S. (2013) *Employment Relations in the Shadow of Recession: Findings from the 2011 Workplace Employment Relations Study*, Palgrave Macmillan
- White, M. and Bryson, A. (2013) 'Positive Employee Attitudes: How Much Human Resource Management Do You Need?', *Human Relations*, 66, 3: 385-406
- Wu, N., Hoque, K., Bacon, N. and Llusar, J. C. B. (2015) "High-performance work systems and workplace performance in small, medium-sized and large firms", *Human Resource Management Journal*, 25, 4: 408-423

Table 1: Incidence of PP Among "Matched" Occupations and Employees, cell percentages

	[1]		[2]		[3]		[4]		[5]		[6]	
	All employees		Restricting to employees in occupations where % public sector is > 15% and < 85%		Restricting to employees in occupations where % public sector is > 25% and < 75%		Restricting to employees in occupations where % public sector is > 35% and < 65%		Hard Match		PSM	
	Pub	Pri	Pub	Pri	Pub	Pri	Pub	Pri	Pub	Pri	Pub	Pri
<i>Incidence:</i>												
Individual PP	5	15	4	15	5	12	5	10	6	11	6	10
Team PP	2	9	2	9	2	7	4	7	2	5	2	4
Workplace/Org PP	1	14	2	14	3	13	4	12	1	8	1	7
Any PP	7	27	7	26	9	23	10	22	8	18	8	16
<i>Difference (Private – Public):</i>												
Individual PP		10***		11***		7***		5*		5***		3***
Team PP		7***		7***		5***		3		3***		3***
Workplace/Org PP		13***		12***		10***		8***		7***		6***
Any PP		20***		19***		14***		12***		10***		8***
<i>N</i>	8,048	13,137	3,943	5,954	2,158	2,579	1,223	1,029	4,273	4,273	7,584	11,779

<i>N unit groups</i>	280	348	127	129	71	71	36	36	265	265	85	85
<i>Mean of SOC_pub%</i>	59	18	46	31	47	38	50	48	39	39	39	39

Notes: (1) In comparisons 2-4, we omit small numbers of employees belonging to occupations for which the QLFS estimate for SOC_pub% has a coefficient of variation of 20% or more (the publication threshold for QLFS estimates). In comparison 2, we omit 65 employees from 11 occupations. In comparison 3, we omit 18 employees from 5 occupations. In comparison 4, we omit 6 employees from 2 occupations.

(2) The Hotelling t-test indicated that, after matching, we could not reject the hypothesis that there was no significant difference in the mean values for the matching covariates between public and private sector employees ($p > 0.996$). The pseudo-rsquared for the probit estimating public sector employment was 0.25 in the unmatched sample, falling to 0.00 in the matched sample, indicating that the propensity score reweighting balanced the covariates between public and private sector employees such that they were observationally equivalent.

(3) Key to statistical significance: * $p < 0.10$, && $p < 0.05$, *** $p < 0.01$.

Base: all employees in workplaces with five or more employees.

Table 2: Effects of PP on intrinsic job satisfaction and organisational commitment in the public and private sectors

	Employees receiving any performance pay	Employees receiving fixed pay only	Difference	Significance (t-stat)
Private sector				
<i>Intrinsic job satisfaction</i>				
Unmatched	.097	.026	.071	3.33
Matched	.097	.043	.054	2.31
<i>Organisational commitment</i>				
Unmatched	.138	.017	.121	5.60
Matched	.137	.057	.080	3.40
Public sector				
<i>Intrinsic job satisfaction</i>				
Unmatched	-.238	-.041	-.197	4.40
Matched	-.238	-.233	-.005	0.10
<i>Organisational commitment</i>				
Unmatched	-.273	-.035	-.238	5.37
Matched	-.270	-.240	-.030	0.62

Notes:

(1) In both sectors the Hotelling t-test indicated that, after matching, we could not reject the hypothesis that there was no significant difference in the mean values for the matching covariates between PP and fixed pay employees. In the private sector the pseudo-rsquared for the probit estimating receipt of PP was 0.12 in the unmatched sample, falling to 0.00 in the matched sample, indicating that the propensity score reweighting balanced the covariates between PP and fixed pay employees employees such that they were observationally equivalent. The same was true for the public sector where the pseudo-rsquared was 0.12 in the unmatched sample and 0.01 in the matched sample.

Base: all employees in workplaces with five or more employees.

Table 3: Effects of PP on Workplace Performance in the Public Sector

	Workplaces with any performance pay	Workplaces with no performance pay	Difference	Significance (t-stat)
<i>Additive Workplace Performance Scale</i>				
Unmatched	4.714	5.008	-.295	1.93
Matched, ATT	4.723	5.053	-.331	1.80
Matched, ATNT	4.569	5.029	-.459	-

Notes:

(1) Although the Hotelling t-test indicated that significant differences between the matched and comparator samples remained in the mean scores for matching the pseudo-rsquared for the probit estimating receipt of PP fell from 0.152 ($p > \chi^2 0.000$) in the unmatched sample to a statistically non-significant 0.006 ($p > \chi^2 1.000$) in the matched sample, indicating that the propensity score reweighting balanced the covariates between PP and fixed pay workplaces such that they were observationally equivalent.

Base: all public sector workplaces with five or more employees