Levelling Down: The Distributional Consequences of Public Pay Caps

Matthew Nibloe*

August 6, 2025

Abstract

This paper studies how progressive public sector pay caps shape the overall wage distribution, conditional inequalities and workforce composition. Between 2010 and 2017, nominal wage growth in the UK public sector was capped at 1% for all workers earning above £21,000—the median annual salary in 2010. Utilising administrative data, we show that the cap reduced the return to observable characteristics in the public sector, coinciding with an 80% reduction in the public–private sector wage premium. Our counterfactuals suggest that the *overall* hourly P90–P50 ratio would have been 3.5 percentage points higher had public sector returns to characteristics moved in line with the private sector. Since public workers are disproportionately female and concentrated in the North of Britain, the policy substantially increased both the gender pay gap and the North–South divide. Over the same period, the public sector experienced relative declines in the share of degree-educated workers, 26-55 year-old prime-age workers, and high-skilled managers, further suppressing wage growth.

Key words: Austerity, Public Sector, Pay Caps, Inequality, Brain Drain

JEL: J24, J31, J45

*University College London, RFBerlin and IFS. Email: matthew.nibloe.19@ucl.ac.uk

I would like to offer special thanks to Christian Dustmann and Uta Schönberg for their ongoing support and supervision throughout this project. I am also grateful for useful comments from Brian Bell, Clara von Bismarck-Osten, Richard Blundell, Bernd Fitzenberger, Ararat Gocmen, David Green, Hyejin Ku, Ekaterine Khvitia, Eric Klemm, Thomas Lazarowicz, Stephen Machin, Sukrit Puri, Imran Rasul, numerous seminar participants at AASLE (2024), the RFBerlin-CReAM Annual Conference (2023 & 2024), and the Stone Centre Symposium at the British Academy (2023).

1 Introduction

The public sector plays a central role in developed economies by providing essential services with positive externalities—such as education, healthcare, and policing. Across the OECD, it accounts for approximately 18% of total employment, making competitive public wage policy a key driver of productivity and output (Britton & Propper, 2016; Burgess et al., 2017; Fenizia, 2022; OECD, 2025). However, in response to the global financial crisis, governments of over 25 IMF high-income countries implemented public pay restraint policies—including wage freezes and caps—as part of broader fiscal consolidation efforts (Ortiz et al., 2015). Many of these policies remained in place considerably after the financial crisis, and the long-run effects have been associated with industrial action (Cribb & O'Brien, 2024; Michael & Christofides, 2020). Despite their prevalence, comparatively little is known about the distributional and compositional consequences of such policies. Instead, research focuses on within-sector inequality or average public-private pay gaps (Fenizia et al., 2024; Michael & Christofides, 2020).

This paper addresses that gap by examining a well-defined UK policy: the public sector pay freeze from 2010 to 2017, which applied to all public employees earning above £21,000 (approximately $32,000 in 2021 USD). Introduced as part of a broader austerity programme, the policy aimed to reduce government spending from 45% to 38% of national income and eliminate a 10% budget deficit (Emmerson, 2017). Its duration, scope, and clear design make the UK an ideal case study, with approximately half of all public sector workers—over 10% of total employment—being affected.

Over the policy period, real hourly wage inequality in the UK—as measured by the P90–P10 ratio—fell by over 9%, reversing a 30% rise between 1980 and 2010 (Bell et al., 2022; Giupponi & Machin, 2024)[1]. The decline was driven by falling wages at the top and rising wages at the bottom. Most explanations emphasise the latter—such as the National Living Wage (Dube, 2019; Giupponi & Machin, 2022) and rising demand for low-skilled work (Giupponi & Machin, 2024)—but cannot account for the 15% decline

[1] Calculated from the sample of workers aged 16 and over in the Annual Survey of Hours and Earnings.

at the 90th percentile. We revisit this contraction by focusing on the P90–P50 gap and assess the extent to which public pay restraint explains the fall in top-end inequality.

First, we document a 4.8 percentage point reduction in the *public sector* P90–P50 ratio between 2010 and 2018. We then simulate counterfactual wage growth along the public sector distribution, assuming that the return to observable characteristics in the public sector mirrored those in the private sector, using the decomposition method of DiNardo et al. (1996). Under this counterfactual scenario, the public sector P90–P50 ratio would have increased by 2.1 percentage points.

Treating private sector wages as observed, we extend our analysis to show that the *overall* P90–P50 ratio in the economy would have been 3.5 percentage points higher in 2018 under the same counterfactual scenario. The pronounced effect on overall inequality arises because public sector workers are disproportionately concentrated at the upper end of the overall wage distribution. This underscores the importance of public sector pay policies in shaping broader income inequality trends during this period.

Our second set of findings examines conditional inequalities. With nearly 70% of public sector workers being women, the pay caps disproportionately affected female wages. Our counterfactual analysis indicates that changes to public sector pay policy led to a 14.5% increase in the gender pay gap in real terms, relative to the counterfactual scenario. Similarly, the public sector accounts for a larger share of employment in the North of Britain (including Scotland and Wales), a share that has grown since 2010. As a result, workers in these regions were more exposed to public pay austerity, leading to a 3.9% increase in the wage gap between the North and the rest of Britain.

Turning to changes in composition, we analyse how the pay caps, in conjunction with concurrent budget cuts—and the resulting decline in the public sector share of employment from 32% in 2010 to 24% in 2018—altered the observable characteristics of public workers and contributed to changes in wage inequality. After 2010, the public sector retention rate—measured as the share of public sector workers who remain in the sector the following year—fell by 8.1 percentage points, reversing the pattern observed in the previous seven

years, during which public sector retention consistently exceeded that of the private sector. Among workers exiting the public sector, the majority transitioned directly into private sector employment rather than into unemployment.

We show that the outflow rate increased along the wage distribution relative to the private sector, and that workers with general skills that are marketable in the private sector—such as public managers—were disproportionately likely to separate compared to those with specialist skills, such as medical and teaching professionals. Furthermore, after accounting for changes in occupational composition, we find that shifts in observable worker characteristics—including age, region, and sex—contributed to a 2.8% reduction in median public sector wage growth relative to the private sector, although the distributional consequences of these compositional changes were minor.

Our paper builds on Bradley et al. (2017), Michael and Christofides (2020), and Di Porto et al. (2023), who examine public–private wage differentials from complementary angles. Bradley et al. (2017) develop a structural model to assess the allocative efficiency of sectoral wage gaps in the UK, simulating the effects of hypothetical wage reforms. In contrast, we evaluate the realised impact of the 2010 UK public sector pay cap using reduced-form methods and administrative microdata, tracing actual changes in wages and composition under a binding institutional constraint. Michael and Christofides (2020) document top-end public sector wage compression across Europe during the austerity period, but do not quantify effects on overall inequality or specific subgroups.[2] Di Porto et al. (2023) study a similar reform in Italy, identifying causal effects of pay caps on high-earning public managers. While their design demonstrates that caps prompt high-performers to exit, their sample is small and selective.[3] Our descriptive analysis, by contrast, covers a far larger and more representative population.

To our knowledge, this is the first paper to open the black box on the distributional consequences of austerity driven by expenditure reductions in the UK. Existing research on the distributional effects of austerity has largely focused on tax changes (Cribb et

[2]See also Ortiz et al. (2015) for a comprehensive list of public pay policies across countries.
[3]Di Porto et al. (2023) study a pay cap for public sector managers that binds at €230,000.

al., 2023), without disaggregating impacts across the public and private sectors. Studies that do examine expenditure cuts—such as Fetzer (2019), Villa (2024), and Facchetti (2024)—tend to focus on specific services or sectors and estimate causal effects using local variation. While these contributions are valuable, they are more limited in scope. In contrast, our counterfactual approach allows us to examine economy-wide changes in the wage distribution—across sectors and subgroups—albeit under stronger partial equilibrium assumptions.

The results in this paper also contribute to the literature on institutional explanations of wage inequality (see, for example, DiNardo et al. (1996); Lemieux (2006)). Most existing work has focused on factors that increase wages at the bottom of the distribution, such as the minimum wage (Dickens & Manning, 2004; Dube, 2019; Giupponi & Machin, 2022) or unionisation (Blanchflower & Bryson, 2004; Dustmann et al., 2009). In contrast, our paper highlights the role of institutional forces in shaping wage dynamics at the top of the distribution.

The paper proceeds as follows: Section 2 describes our administrative data. Section 3 outlines changes to the public sector wage structure before and after 2010, along with the relevant austerity policies. Section 4 details our counterfactual methodology and presents summary statistics. Section 5 presents our main results on overall inequality, the gender pay gap, and the North–South divide. Section 6 concludes.

2 Data

The primary dataset for this project is the Annual Survey of Hours and Earnings (ASHE). Our administrative data span 1997 to 2019 and links workers to their employers at the individual level. The dataset comprises a 1% random sample of employees, selected based on the last two digits of their National Insurance number—the UK equivalent of the US Social Security number. ASHE is the foremost source of earnings data in the UK and serves as the basis for the Office for National Statistics' (ONS) official earnings statistics. It has also been widely used in both economic and policy research.

Individual tax records—collected via His Majesty's Revenue and Customs (HMRC) Pay As You Earn (PAYE) system—are linked to firms registered on the Inter-Departmental Business Register (IDBR). Data are collected directly from employers, which minimises attenuation bias typically associated with self-reported surveys, such as the Labour Force Survey (LFS). Participation is legally mandated, and the response rate is approximately 80%, corresponding to around 180,000 individuals annually.

Despite being a theoretically random sample, certain individuals are not observed. Firms with annual turnover below £85,000 ($106,000) are not required to register with the IDBR, implying that many self-employed sole traders are omitted. Bell et al. (2022) note that these workers account for approximately 15% of total employment. Moreover, workers who do not pay National Insurance—those earning up to approximately £120 ($151) per week—are not registered in the PAYE system and therefore cannot be linked to a firm.[4] This is equivalent to a minimum wage worker employed for 15 hours per week, representing a very small proportion of the population, primarily composed of women in the private sector.

A public sector worker is defined as an employee in central or local government, or in public corporations.[5] In the UK, this includes occupations such as teachers and civil servants. Unlike in the US, the public sector also includes the majority of healthcare professionals, who are employed by the National Health Service (NHS). By this definition, general practitioners are not considered public workers. Although they are publicly funded, they are self-employed and operate their own practices. Additionally, public universities and their professors are classified as private sector employees, as these institutions have wage-setting power. Across our sample, public workers account for 28.5% of total employment, suggesting that any change to public sector pay has implications for the overall wage distribution.

Our primary outcome variable is gross hourly wage, calculated as weekly earnings (in-

[4]National Insurance is the statutory tax that all workers pay to be eligible for a state pension.

[5]This follows the formal definition adopted by the Office for National Statistics, the UK's official statistical body.

cluding overtime) divided by the number of hours worked in the reference week.[6] Additional variables include occupation, sex, age, local authority region, and union coverage via collective bargaining agreements. While the administrative data provide high-quality earnings information, they do not capture labour market states beyond employment in the public or private sector. To address this, we supplement our analysis with the harmonised Understanding Society and British Household Panel Surveys (UKLS). The UKLS is a self-reported longitudinal dataset covering approximately 40,000 individuals in the UK and is widely used in the macro-labour literature (Bradley et al., 2017; Sepahsalari & Postel-Vinay, 2023).[7] Crucially, it provides information on employment type (e.g. self-employment) as well as on non-employment states—including unemployment, inactivity, and retirement—that are not observable in the administrative data. This allows us to construct a more complete picture of worker transitions and sectoral flows.

3 Public Sector Wages and Composition

Next, we describe trends in real wages and inequality in the public and private sectors between 1997 and 2018 (Section 3.1). We restrict the distributional analysis to the P90–P50 differential, which is appropriate given that the pay caps primarily affected the top end of the wage distribution. Moreover, the P90 and P50 are robust to the introduction of the living wage and to changes in unemployment, both of which may bias analyses based on lower percentiles.[8] [9] Section 3.2 explores the demographic composition of the public sector and how it evolved after 2010. Finally, Section 3.3 shows that changes in the relative composition of the public sector were predominantly driven by increased outflows at the top end of the wage distribution.

[6]Wages are deflated by the Consumer Price Index (CPI)—the headline measure of inflation in the UK—indexed to 2019.

[7]The UKLS refers to the combination of two surveys: the British Household Panel Survey (BHPS) and Understanding Society (USoC).

[8]The living wage replaced the minimum wage in the UK in 2016, with the goal of raising the wage floor to 60% of median income by 2020. The policy is estimated to affect 7% of individuals aged 25 or over (Dube, 2019).

[9]Between 2010 and 2018, the unemployment rate fell from 7.9% to 4.1%. Blundell et al. (2014) show that the initial unemployment shock was concentrated among young, low-educated male workers, disproportionately affecting the lower percentiles of the wage distribution.

3.1 Trends in Real Wages

Figure 1 shows wage growth from 1997 to 2019 in the public and private sectors at the 90th percentile, normalised to 1997. Wages increased at the upper end of the distribution in both sectors between 1997 and 2010, with public sector wage growth closely tracking that of the private sector. In contrast, wages declined across both sectors between 2011 and 2015, as the UK experienced the prolonged effects of the Financial and Eurozone crises. From 2015 onward, real wages began to recover in the private sector, while public sector wages continued to decline in real terms. Between 2010 and 2018, public sector workers at the 90th percentile experienced 8% lower real wage growth than their private sector counterparts.[10]

We hypothesise that much of the change in the public wage distribution can be attributed to the public sector pay freezes. The policy was introduced in 2010 under Chancellor of the Exchequer George Osborne, who froze the nominal wages of all public sector employees earning above £21,000, with the goal of saving £3.3 billion per year. Workers below the cap threshold were entitled to a minimum raise of £250 per year. For the average public sector worker in 2010—who worked 33.7 hours per week—this placed the cap at £12.00 per hour (2.54 log points), equivalent to the median of the public sector wage distribution.

The nominal pay freeze remained in effect until 2013, when it was loosened to permit 1% annual wage increases. In 2017, the policy was formally repealed, and all public sector workers became entitled to a minimum 2% nominal pay rise.[11] Over the same period, inflation averaged between 1–3%, resulting in persistent real wage declines for workers subject to the cap.

We assume that, in the absence of the pay caps, wage growth above the median—the bite point of the policy—would have mirrored that of the private sector. This assumption is motivated by the historical alignment of public and private sector wages at the 90th

[10]Figure A1 replicates this analysis using weekly wages, weekly wages excluding overtime, and weekly wages including pension contributions as alternative outcome variables. The direction and magnitude of the relative decline in public earnings are consistent across specifications.

[11]The policy was repealed as the government neared its target of balancing the budget and faced increasing union pressure due to the widening gap between public and private sector wage growth.

percentile, both in levels and in trends.[12] Institutionally, public sector wages are set based on the recommendations of independent government pay review bodies, which explicitly consider private sector wage trends in their decisions. Prior to 2010, these recommendations were routinely adopted by the government. Between 2010 and 2017, however, they were overridden in favour of the pay cap.[13]

Therefore, we assume that between 2010 and 2015, the pay caps did not bind in relative terms, as public and private sector wages moved in parallel. From 2015 onward, however, the caps began to have a real effect: the public–private wage differential widened, and the relative value of outside private sector offers increased. We interpret our counterfactual as a lower bound, since the relative decline in public sector wages over this period likely heightened competition for private sector employment, exerting downward pressure on private sector wages.[14]

3.1.1 Wage Growth Along the Distribution

To illustrate the impact of the pay caps, Figure 2 plots wage growth along the distribution relative to 2004. The left panel shows the public sector; the right panel shows the private sector. In the public sector, wage growth was approximately uniform across the distribution. By contrast, private sector wage growth was mildly regressive, with higher earners experiencing slower growth. Wage growth remained positive in the public sector between 2004 and 2010, as the government continued to honour pre-arranged pay agreements following the 2007 recession. In the private sector, however, wages declined due to the immediate market response to the crisis.

The red line in Figure 2 represents wage growth between 2004 and 2018. In the public sector, it falls below the pre-2010 trend (black line) at all percentiles, indicating that wage growth was negative across the distribution during the pay cap period. The decline

[12] See Table 1 for evidence on pre-2010 public-private equivalence at the 90th percentile.

[13] For example, the 2023 NHS Pay Review Body report (paragraph 3.215) compares the wages of health and social care workers in the public and private sectors. Similar comparisons are featured in each annual report.

[14] Gregg et al. (2014) document that wages became increasingly sensitive to unemployment after 2010, suggesting that workers were more willing to accept lower wage offers rather than risk unemployment, in contrast to earlier recessionary periods.

was particularly pronounced above the cap threshold. In contrast, private sector wages remained largely flat across the distribution over the same period, with a modest increase at the bottom.

3.2 Characteristics of Public Sector Workers

The overall inequality effect of changes to public sector wages and composition depends on where public sector workers are located within the overall wage distribution. This position is partially determined by the observable characteristics of public sector employees. Table 1 presents descriptive statistics for public and private sector workers in 2010 and 2018, respectively.

In 2010, relative to the average private sector worker, public sector employees were 24 percentage points more likely to hold a degree and were, on average, 3.5 years older. Consistent with prior findings (Garibaldi et al., 2021; Postel-Vinay, 2015), these differences suggest that the public sector contains significantly more human capital, on average, than the private sector. This reflects, in part, the occupational structure of the UK public sector, where the state is the monopoly provider of healthcare and education—two high-skill professions in which degree-level education is typically required.

However, the public sector wage distribution is more compressed than that of the private sector. In 2010, the 10th percentile of public sector wages was 22% higher than the corresponding value in the private sector, while at the 90th percentile, the difference was only 8%. This pattern partly reflects the substantial returns available at the top of the private sector wage distribution, which are largely absent in the public sector (Bell & Van Reenen, 2014).

3.2.1 Changes over Time

As part of the government's fiscal retrenchment, public sector employers were required to reduce vacancy postings and refrain from renewing temporary contracts (Andrews, 2015; ONS, 2019). Moreover, as the pay caps began to bind, we hypothesise that public

sector workers voluntarily exited the sector due to the declining relative value of public employment. Together, these factors are associated with an 8 percentage point decline in the public sector share of total employment between 2010 and 2018, as shown in Figure 3.

Columns 3 and 6 of Table 1 summarise how these factors changed the composition of the public sector by reporting the percentage change in observable worker characteristics between 2010 and 2018. Over this period, the average level of education increased in both sectors. However, the growth in the share of degree-educated workers was significantly smaller in the public sector than in the private sector. This suggests that, after 2010, the private sector was more successful in attracting a larger share of skilled labour.

Second, the relative share of men in the public sector declined. A large literature suggests that men in the public sector tend to earn a wage premium over women (Garibaldi et al., 2021; Singleton, 2019), driven in part by differences in labour supply elasticities and employer discrimination (Andrew et al., 2024; Mulligan & Rubinstein, 2008). As a result, the reduction in the male share likely shifted the public sector wage distribution leftward by reducing the average wage level.

Third, the share of public sector workers located in the North of Britain increased.[15] As noted by McCann (2020) and Overman and Xu (2024), average earnings and education levels are significantly higher in the South. Accordingly, the shift in public employment toward the North likely contributed to a decline in average public sector wages.

Figure 4 shows how the occupational distribution in the public sector changed relative to 1997 by decomposing workers into three categories: low-skilled occupations, teachers and medical professionals, and remaining high-skilled occupations.[16] Teachers and medical professionals are separated from other high-skilled occupations because they are almost exclusively employed in the public sector, which limits between-sector mobility and reduces the availability of outside employment options.

[15]The North is defined using government office regions north of the Midlands, and includes Scotland and Wales.

[16]We define an occupation as low-skilled if the share of degree-educated workers in that occupation is below the sample mean.

The results suggest that the public sector was able to exercise monopsony power in these occupations, reducing wages without a corresponding change in occupational composition. The decline in the share of high-skilled public sector workers is driven primarily by two occupations. First, the share of managers and senior officials fell from 4.0% to 2.7% of total public employment between 2011 and 2018. Second, the share of health and social care associate professionals declined from 3.1% to 2.8% over the same period.

The decline in the share of highly skilled public sector workers suggests a brain drain from the sector, alongside occupational consolidation around core public services such as education and healthcare.

3.2.2 Public Sector Conditional Premium

Given that public sector workers differ substantially in observable characteristics from their private sector counterparts, comparing raw wage percentiles does not yield a like-for-like comparison. Figure 5a shows that even after conditioning on these characteristics, the pay caps reduced the relative value of public sector job offers—effectively increasing the opportunity cost of remaining in public employment.

The black line in the figure plots the mean difference in real wages between public and private sector workers, referred to as the public sector wage premium. This premium rose steadily from 17% in 2005 to 21% in 2010, before declining to 16% by 2018.

The red line in Figure 5 adjusts for selection on observables by accounting for compositional effects. Specifically, it reflects the estimated mean wage differential after conditioning on age (and age squared), sex, union status, their interactions, and individual, regional, and occupational fixed effects. This provides a closer approximation of the wage that a public sector worker could earn in the private sector. Formally, this involves estimating separate wage equations by sector and comparing the mean residuals across time.[17]

[17] This approach is equivalent to estimating the equation:

$$\ln(w_{it}) = \alpha_t \cdot \text{public}_{it} + \beta_t^s X_{it} + \phi_i + \theta_t + \varepsilon_{it},$$

where α_t is the coefficient on an indicator for public sector employment, w_{it} is the hourly wage, X_{it} is a

The conditional premium follows a similar trend to the unadjusted measure, though it is lower in magnitude. Between 2010 and 2018, it fell from 9.4% to 1.9%, effectively eliminating the wage-based incentive to work in the public sector.[18]

Figure 5b examines the decline in the conditional public sector premium across the wage distribution between 2010 and 2018. The reduction was evident throughout the distribution but was particularly pronounced above the median—where the pay caps applied—reinforcing the pattern observed in Figure 2.

3.2.3 Relative Position of Public Workers in the Distribution

Given the higher average level of observable human capital in the public sector relative to the private sector, we would expect public sector workers to be concentrated toward the top of the overall wage distribution. This pattern is illustrated in Figure 6, which plots the public sector share of employment along the log hourly wage distribution in 2010 and 2018, using 0.1 log-point intervals. In 2010, the public sector share peaked at 48.5% at 3.0 log points (approximately the 75th percentile). In contrast, at 2.0 log points—around the 10th percentile—the share was only 16.4%.

However, following the decline in the public sector wage premium after 2010 and the exit of high-skilled workers, we would expect the relative rank of public sector employees within the overall wage distribution to fall—particularly at the top end. The red line in Figure 6 illustrates this pattern. At every point along the distribution, the share of public sector workers was lower in 2018 than in 2010, reflecting the overall decline in public sector employment.

Importantly, the drop was not uniform. Above the pay cap threshold, the public share fell by 9.5 percentage points, compared to just 3.7 percentage points below the threshold. This implies that public sector workers became increasingly underrepresented at the top

vector of individual characteristics (including region and occupation fixed effects), and ϕ_i and θ_t denote individual and time fixed effects, respectively.

[18]This result is consistent with findings from Boileau et al. (2022), who estimate a comparable decline using Labour Force Survey data. Our analysis relies on high-quality administrative records, though it does not include educational attainment directly; this is partially proxied by occupation fixed effects.

of the wage distribution.

3.3 Inflow and Outflow

The decline in the public sector share of employment may reflect either an increase in outflows from the public sector or a reduction in inflows. This subsection examines which of these channels contributed most to the compositional changes described above, drawing on data from both UKLS and ASHE.

Table 2 presents a transition matrix of annual labour market flows. Panel A shows the average annual transition probabilities between labour market states from 2004 to 2009. Panel B reports the same for the post-reform period, 2011 to 2018. Panel C displays the percentage point change in transition probabilities between the two periods.[19]

Focusing on the public sector in Panel A, 88% of workers remained in public employment in the subsequent year during the pre-2010 period. Public sector retention—as measured by the proportion of workers staying in the same state year to year—was higher than for any other labour market state. Among those who exited the public sector, most transitioned into private sector employment, followed by inactivity. Compared to private sector employees, public sector workers were substantially more likely to remain in employment overall. These patterns underscore the high degree of job security associated with public sector employment prior to 2010.

After 2010, outflows from the public sector increased by 8.1 percentage points, while inflows declined by only 0.3 percentage points. Approximately 85% of the increased outflow consisted of workers transitioning to the private sector, with most of the remainder moving into retirement. As a result, public sector retention fell to 79.9%—below the private sector retention rate of 83.0%.

These figures suggest that post-2010 changes in the composition of the public sector wage

[19]We omit 2010 from the analysis because it marks the transition from the British Household Panel Survey (BHPS) to Understanding Society (USoC), which introduced a change in the sampling structure. To avoid bias, we also reweight the data to ensure a balanced cross-section and to adjust for composition effects related to ageing.

distribution were primarily driven by increased outflows, rather than reduced inflows.

3.3.1 Outflow Along the Wage Distribution

Figure 7 compares the percentage point change in the probability of exiting a given sector across the wage distribution before and after 2010, using five equally spaced quintiles.[20] Outflow increased at all points in both the public and private sectors. Below the median, the change in outflow was similar across sectors—around 7 percentage points. However, above the median, and particularly in the top quintile, the divergence is notable: public sector outflow increased by 7.0 percentage points, compared to just 3.0 percentage points in the private sector. This pattern suggests that the contraction of the public sector after 2010 was primarily driven by increased outflows, especially among higher earners.[21]

4 Counterfactual Methodology

Can the changes in wage inequality since 2010 be attributed to public sector pay restraint and expenditure cuts? In this section, we formalise the distributional impact of austerity by developing a log-linear model of wages (Section 4.1), which allows us to decompose changes in the wage distribution into price and composition components.[22]

In Section 4.2, we apply the reweighting procedure of DiNardo et al. (1996) to estimate the wage distribution that would have prevailed in 2018 if observable characteristics and their returns had remained at 2010 levels. We refer to this as the *characteristic effect*.

We then incorporate an Oaxaca–Blinder-style decomposition to isolate the *price effect*, capturing changes in returns to characteristics over time, holding composition fixed. Next, this framework is extended to estimate the *relative price effect*—the counterfactual public

[20] We define the pre-2010 period as pooled data from 2009–2010 and the post-2010 period as 2017–2018. Pooling is necessary due to the relatively low frequency of transitions, particularly when disaggregated along the wage distribution.

[21] Appendix Figure A2 replicates the analysis for inflows and finds that changes were broadly similar across sectors up to the top quintile. At the very top, the decline in inflows was smaller in the public sector than in the private sector. However, these differences are modest relative to the outflow patterns, reinforcing the conclusion that outflows were the dominant driver of compositional change.

[22] For reference, Table A1 summarises the counterfactual effects identified throughout the analysis.

sector wage distribution that would have emerged if public sector prices had evolved in line with the private sector, under constant composition.

Finally, in Section 4.3, we map these effects onto the overall wage distribution and derive the *share effect*, which isolates the contribution of changes in the public sector employment share across the wage distribution.

4.1 Decomposing Densities

Log wages, w_{it}, are generated according to Equation 1. Each worker possesses a vector of observable characteristics, x_{it}, which includes biological sex, hours worked, age (and age squared), union status, their interactions, alongside occupation and region fixed effects.[23]

$$w_{it} = x_{it}\beta_t^s + \varepsilon_{it} \tag{1}$$

Observable characteristics generate labour market returns via a sector-specific price vector, β_t^s, where $s \in \{\text{public, private}\}$. This captures both the relative compression of public sector wages and variation in returns across sectors. Prices vary over time to reflect the observed wage decline between 2010 and 2018. Wages also include an idiosyncratic component, ε_{it}, capturing unobservable factors—such as those related to Brexit or skill-biased technical change—that influence the supply and demand for specific skills (Juhn et al., 1993).

An Oaxaca–Blinder decomposition allows us to describe how the wage distribution evolves within our model. Suppose we observe two periods, 2010 and 2018. The change in the

[23]We use two-digit occupation controls to ensure each occupation contains both public and private sector workers, preserving comparability across sectors.

wage density function in sector s can then be decomposed as:

$$\underbrace{\Delta f_s(w \mid \beta_s, x)}_{\text{Overall change}} = f_s^{18}(w \mid \beta_s^{18}, x^{18}) - f_s^{10}(w \mid \beta_s^{10}, x^{10})$$

$$= \underbrace{f_s^{CF}(w \mid \beta_s^{18}, x^{10}) - f_s^{10}(w \mid \beta_s^{10}, x^{10})}_{\text{Price effect } = \Delta f_s(\beta)} + \underbrace{f_s^{18}(w \mid \beta_s^{18}, x^{18}) - f_s^{CF}(w \mid \beta_s^{18}, x^{10})}_{\text{Characteristic effect } = \Delta f_s(x)} \tag{2}$$

The first term represents the price effect, capturing changes in returns to observable characteristics over time. In the public sector, where wages are largely set by government policy, this reflects institutional changes such as the introduction of the National Living Wage in 2016 and public sector pay caps. The price effect thus isolates shifts in wage-setting structures while holding worker characteristics constant.

To interpret this effect as the impact of the pay caps, we make two assumptions. First, the caps had no spillover effects below the threshold. Second, the shape of the wage distribution above the threshold depends only on the magnitude of the caps. These assumptions parallel those used to identify minimum wage effects in DiNardo et al. (1996).

The second term is the characteristic effect, capturing changes in the wage distribution driven by selective inflows and outflows that alter the composition of observable characteristics. This aggregate effect reflects the combined impact of hiring freezes, job destruction, and voluntary exits in response to pay restraint. While we do not isolate these channels individually, changes in workforce composition offer insight into how austerity affected the observable characteristics of public sector employment, which may correlate with output.

4.2 Calculating Characteristic and Price Effects

To quantify how changes in worker characteristics influenced the wage distribution, we ask how the distribution would have evolved if the composition of the workforce had remained fixed at its 2010 levels. The DiNardo et al. (1996) reweighting approach allows us to isolate this characteristic effect by decomposing the wage distribution at any time

t into two components: a price component, $h_t(w \mid x)$, which reflects the conditional wage distribution given characteristics, and a characteristic component, $g_t(x)$, which captures the distribution of those characteristics. Formally, the observed distribution of wages in 2018 is given by:

$$f_s^{18}(w \mid \beta_s^{18}, x^{18}) = \int h_s^{18}(w \mid \beta_s^{18}, x) \cdot g_s^{18}(x)\, dx \tag{3}$$

To identify the characteristic effect, we recover the counterfactual wage distribution $f_s^{CF}(w \mid \beta_s^{18}, x^{10})$ from Equation 2, representing the distribution that would have prevailed in 2018 had observable characteristics remained at their 2010 levels, assuming prices followed their observed 2018 path. The key insight of DiNardo et al. (1996) is that this counterfactual can be obtained by reweighting the observed 2018 wage distribution in Equation 3 to reflect the 2010 characteristic distribution:

$$f_s^{CF}(w \mid \beta_s^{18}, x^{10}) = \int h_s^{18}(w \mid \beta_s^{18}, x) \cdot g_s^{18}(x) \cdot \phi_s(x)\, dx \tag{4}$$

Here, $\phi_s(x)$ denotes the reweighting function, interpreted as an inverse probability weight. If a given set of observable characteristics x was more common in 2010 than in 2018, then the corresponding 2018 observations are upweighted by a value greater than one. This procedure ensures that the reweighted distribution of characteristics in 2018 matches the 2010 distribution exactly. Formally, the weight is defined as:

$$\phi_s(x) = \frac{g_s^{10}(x)}{g_s^{18}(x)} = \frac{\Pr(t = 10)}{\Pr(t = 18)} \cdot \frac{\Pr(t = 10 \mid x)}{\Pr(t = 18 \mid x)} \tag{5}$$

Each of the components in the reweighting function can be estimated using binary choice models, such as logit or probit regressions predicting period membership based on observable characteristics. The counterfactual wage density defined in Equation 4, along with the observed densities in 2010 and 2018, can be estimated non-parametrically using kernel density methods. Given estimates of each of these three densities, we can apply

the Oaxaca–Blinder identity from Equation 2 to recover both the price and characteristic effects.

4.2.1 Isolating the Effect of Individual Characteristics

One of the most notable structural shifts in the public sector over this period is the growing concentration of employment in teaching and medical professions. The DiNardo et al. (1996) framework offers a parsimonious way to isolate this effect by modifying the reweighting function in Equation 5. Specifically, we use the alternative weight in Equation 6 to construct a counterfactual 2018 wage distribution in which the occupational share of teachers and medical professionals is held at its 2010 level, while returns to characteristics and all other covariates follow their observed 2018 values.

Let o be a binary indicator equal to 1 if the individual is a teacher or medical professional, and let x^{-o} denote the covariate vector excluding o. The reweighting function is:

$$\theta_s(o \mid x_{-o}) = o \cdot \frac{\Pr(o = 1 \mid x_{-o}, t = 10)}{\Pr(o = 1 \mid x_{-o}, t = 18)} + (1 - o) \cdot \frac{\Pr(o = 0 \mid x_{-o}, t = 10)}{\Pr(o = 0 \mid x_{-o}, t = 18)} \tag{6}$$

4.2.2 Capturing Relative Price Changes

The price effect measures the impact of pay caps on public sector wages relative to a 2010 benchmark. As noted above, it is unlikely that wages would have remained flat in the absence of caps. One reason is the significant decline in private sector wages between 2010 and 2014, following the financial crisis. Given the historical comovement between public and private sector wages, we treat the path of private sector returns to observable characteristics from 2010 to 2018 as the counterfactual path for public returns in the absence of austerity. We refer to this as the *relative price effect*, which can be recovered

by decomposing the price effect from Equation 2 as follows:

$$\Delta f_{\text{pub}}(\beta) = \underbrace{f^{CF}_{\text{pub}}(w \mid \beta^{18}_{\text{pub}}, x^{10}) - f^{CF}_{\text{priv}}(w \mid \beta^{18}_{\text{priv}}, x^{10})}_{\text{Relative price effect}} + \underbrace{f^{CF}_{\text{priv}}(w \mid \beta^{18}_{\text{priv}}, x^{10}) - f^{10}_{\text{pub}}(w \mid \beta^{10}_{\text{pub}}, x^{10})}_{\text{Residual price effect}} \tag{7}$$

Each component of Equation 7 is either directly observed or can be estimated using the DiNardo et al. (1996) reweighting procedure outlined above.

4.3 Mapping Public Changes to the Overall Wage Distribution

To map the previously defined effects onto the overall wage distribution, we express the aggregate distribution as a wage-dependent weighted average of the public and private sector distributions. Since public sector employment shares vary across the wage distribution, the sector weights must also vary with wage. Let $\delta(w \mid x)$ denote the share of public sector workers with characteristics x at wage level w; then the overall distribution is given by:

$$f^t_{\text{all}}(w \mid \beta^t_{\text{all}}, x^t) = \delta(w|x) \cdot f^t_{\text{pub}}(w \mid \beta^t_{\text{pub}}, x^t) + [1 - \delta(w|x)] \cdot f^t_{\text{priv}}(w \mid \beta^t_{\text{priv}}, x^t) \tag{8}$$

Assuming that sectoral employment shares remain constant over time, the change in the overall wage density between 2010 and 2018 can be expressed as a weighted average of the changes in the private and public sector wage densities:

$$\Delta f_{\text{all}}(w \mid \beta_{\text{all}}, x) = \delta(w|x) \cdot \Delta f_{\text{pub}}(w \mid \beta_{\text{pub}}, x) + [1 - \delta(w|x)] \cdot \Delta f_{\text{priv}}(w \mid \beta_{\text{priv}}, x) \tag{9}$$

Combining Equations 2, 7 and 9 allows us to identify the change to the wage distribution

due to each of the previously defined counterfactual scenarios:

$$\begin{aligned}\Delta f_{\text{all}}(w \mid \beta_{\text{all}}, x) = \delta(w|x) \cdot \overbrace{\Bigg[\underbrace{\Delta f^{\text{rel}}_{\text{pub}}(\beta)}_{\text{Relative price}} + \underbrace{\Delta f^{\text{resid}}_{\text{pub}}(\beta)}_{\text{Residual price}} \Bigg]}^{\Delta f_{pub}(\beta)} \\ + \delta(w|x) \cdot \overbrace{\Bigg[\underbrace{\Delta f^{o}_{\text{pub}}(x)}_{\text{Occupations}} + \underbrace{\Delta f^{-o}_{\text{pub}}(x)}_{\text{Other characteristics}} \Bigg]}^{\Delta f_{pub}(x)} \\ + [1 - \delta(w|x)] \cdot \overbrace{\Delta f_{\text{priv}}(w \mid \beta_{\text{priv}}, x)}^{\text{Private sector distributional shifts}}\end{aligned} \tag{10}$$

To estimate the effects of public sector austerity on conditional inequalities—such as gender or regional disparities—we repeat the decomposition separately for men and women, and for workers in the North and South, estimating sector-specific wage densities within each subgroup.

This analysis relies on a partial equilibrium assumption: private sector wages and composition do not respond to changes in public sector pay or workforce characteristics. We justify this assumption on three grounds. First, the public sector labour force is about one-third the size of the private sector, so even a large outflow of public workers to the private sector would generate only a modest increase in private sector labour supply. Second, as discussed later in Section 5.1 (Figure 8b), compositional change contributed little to private sector wage growth between 2010 and 2018. Third, Bradley et al. (2017) find that private sector wage offers respond only marginally to public pay cuts: a cap reducing the public sector wage bill by 10%—comparable to the fall in real public wages over the period—would change private offers by less than 1%. This claim is supported descriptively in Figures 2 and 8a, which show a clear downward shift in real wage growth in the public sector between 2010 and 2018 around the pay cap threshold. In contrast, there is no significant change in private sector wage growth. Taken together, this evidence sug-

gests that any spillover effects into the private sector are likely to be small in magnitude and not substantial enough to materially alter our conclusions about the distributional impact of public pay restraint.

4.3.1 Accounting for Changes to the Public Sector Share

To account for the decline in the public sector employment share between 2010 and 2018, we relax the assumption that $\delta(w \mid x)$ is fixed over time. We estimate $\delta(w \mid x)_t$ for each year using a binary choice model, and use these estimates to calculate $f^t_{\text{all}}(w \mid \beta^t_{\text{all}}, x^t)$ via Equation 8. To construct the change in the wage distribution owing to the change in the public share, we compute the wage distribution twice: once using 2018 weights and once using 2010 weights. The difference between the two yields the share effect.

This approach is analogous to the individual characteristic reweighting method of DiNardo et al. (1996) described in Section 4.2.1, where the characteristic of interest is public sector employment and the counterfactual is calculated using both public and private sector workers.

5 Counterfactual Results

Next, we examine the impact of the austerity measures on inequality between 2010 and 2018 by combining the counterfactual methods outlined above. We begin by assessing the effects of changes in prices and composition across the public and private wage distributions, before turning to their impact on the overall P90–P50 ratio. After demonstrating the progressive effects of the austerity programme along the distribution, we consider their implications for conditional inequalities. In particular, we show that changes in relative public sector prices spilled over into the gender pay gap and substantially widened the North–South wage divide.

5.1 Characteristics, (Relative) Prices and the Public Share

Figure 8a illustrates the price effect in the public and private sectors. The black line represents observed real wage growth along the distribution between 2010 and 2018 in both sectors. As expected, wage growth was negative and progressive across the public wage distribution during this period, whereas private sector wage growth was approximately zero and uniform along the distribution. The grey line indicates the wage growth that would have occurred if prices had remained fixed within each sector. In this counterfactual case, wage growth would have been roughly equivalent below the cap threshold. Above the threshold, however, wage growth would have been significantly higher, reaching 7.7 percentage points at the P90.

Turning to the private sector, we observe that above the cap threshold, wage growth would have been approximately equal to that in the counterfactual case, implying that there was little change in the return to observable characteristics during this period. As discussed earlier, we argue that the wage distribution that would have prevailed if public sector returns to observable characteristics had followed those in the private sector—the relative price effect—provides a more appropriate counterfactual for identifying the impact of the pay caps. The difference between the black and green lines in the public sector panel represents the relative price effect. Between the 50th and 90th percentiles, the price effect and the relative price effect are approximately equal in magnitude. Overall, the counterfactual suggests that the public sector P90 would have been 7.7 percentage points higher in the absence of the pay caps. Rather than a 4.2 percentage point decline in the public P90–P50 ratio over the period, inequality would instead have increased by 2.1 percentage points.

Panel 8b conducts a similar exercise, this time focusing on the impact of changes in worker characteristics on wage growth. In this case, the blue line indicates the wage growth that would have prevailed had characteristics remained fixed. In both sectors, wage growth would have been higher than observed if the composition of public sector workers had not changed. This suggests a general shift away from characteristics positively associated with

wages. According to our prior results, three channels contribute to this trend: employment shifted away from London and the South of England; the male labour share declined; and the share of prime-age workers fell.[24]

The relative composition effects are similar in the public and private sectors, suggesting that the observable characteristics of the two workforces evolved in parallel between 2010 and 2018. However, this comparison does not account for the occupational consolidation around high-paying teaching and medical professionals within the public sector, which likely increased observed wage growth. The grey line adjusts for this by calculating real wage growth in each sector under the assumption that the share of teaching and medical professionals remained fixed. Under this counterfactual, public sector wage growth would have been 3 percentage points lower along the wage distribution, while private sector wage growth remained unchanged. The difference between the grey and blue lines represents the pure impact of changes in the regional, gender, age, and occupational composition on wage growth, net of the consolidation around teaching and medical professionals. Overall, these compositional changes resulted in 2.5% lower wage growth in the public sector relative to the private sector at the median.

To account for the 7 percentage point decline in the public sector share between 2010 and 2018, we employ a variation of the DiNardo et al. (1996) approach to estimate the wage distribution that would have prevailed had the public sector share remained at its 2010 level. Figure 9 presents observed and counterfactual wage growth for the economy under this restriction. Wage growth would have been higher at all points along the wage distribution if the public sector share had remained constant. This reflects the positive and significant wage premium that public sector workers earn relative to their private sector counterparts. However, wage growth would have been only marginally higher above the 80th percentile compared to 0.5 percentage points below. This reflects the relatively low public sector premium at the top of the distribution and the overall compression of public

[24]We are cautious in interpreting the absolute level changes implied by this counterfactual, as the data do not include educational attainment. Our focus is on the relative change between the public and private sectors. We consider this relative difference a lower-bound effect, since the increase in educational attainment was slower in the public sector than in the private sector between 2010 and 2018.

sector earnings. Indeed, the residual public sector premium is approximately zero at the top decile (Boileau et al., 2022). Taken together, these results imply that holding the public sector share constant would have modestly boosted aggregate wage growth by 2018.

5.2 Overall Inequality

We demonstrate the impact of the pay caps on the overall wage distribution by combining our counterfactual public sector wages with observed private sector wages. The left panels of Figure 10 display the counterfactual wage changes at the 90th and 50th percentiles of the distribution, respectively. The green line represents the scenario in which returns to observable characteristics in the public sector followed those in the private sector. In each year, wages at the P90 would have been higher, rising by 4 log points by 2018. However, there is no observable effect at the median. This outcome is expected for two reasons. First, as shown above, public sector returns evolved in line with private sector returns at the median over our sample period. Second, the public sector share is substantially higher at the P90 (47%) than at the median (32%).

The characteristic effect is identified by the difference between the green and blue lines. In both cases, the effect is approximately zero, suggesting that changes in the composition of the public sector played little role in shaping wage inequality. The purple line accounts for changes in the public sector share at the P50 and P90, respectively. As shown in Figure 9, wages at the 50th percentile would have increased marginally had the public sector share remained at 2010 levels, while wages at the 90th percentile would have remained unchanged.

The right-hand panel traces the counterfactual effects through to the P90–P50 wage differential. The results suggest that, by 2018, the overall P90–P50 gap would have increased by 3.5 percentage points rather than falling by 0.6 percentage points, exceeding the 2010 level of inequality by 2.9 percentage points.

5.3 The Gender Pay Gap

An important dimension of inequality is the wage gap between men and women. Between 1997 and 2010, the unconditional gender pay gap—a common measure of wage inequality by sex—fell from 30.1% to 19.9%, at an average rate of 0.7 percentage points per year. As shown in the left panels of Figure 11, this decline was disproportionately driven by faster wage growth among women. However, between 2010 and 2018 the gap narrowed only slightly, reaching 16.7%, or an average reduction of just 0.4 percentage points per year. We now ask whether the public sector pay caps can explain this structural shift in the gender pay gap. To motivate this hypothesis, consider the composition of employment in 2018: women accounted for 70% of public sector employees, compared to only 43% in the private sector. This suggests that women were more likely to be affected by the pay caps and therefore experienced lower wage growth.

The left panels of Figure 11 display observed and counterfactual wages for men and women separately over time. As established above, the decline in public sector returns to observable characteristics relative to those in the private sector was the dominant factor shaping the wage distribution over our sample period. Accordingly, we restrict our analysis to this specific counterfactual.

In the male case, wages would have been 0.4% higher in the absence of the pay caps, whereas for women, wages would have been 2.8% higher. The resulting effect on the gender pay gap is illustrated in the right-hand panel of Figure 11. By 2018, the gap would have been 2.5 percentage points—or 14.5%—lower in the absence of the policy. Moreover, the rate of decline in the gender pay gap would have returned to its pre-2010 pace of 0.7 percentage points per year, suggesting that the pay caps were the source of the structural slowdown in its reduction after 2010.

5.4 Levelling Down and Regional Inequality

Our final analysis explores the regional effect of the pay caps through the lens of the North–South wage divide. Throughout this section, we define the North as any region in

England north of the Midlands, as well as Scotland and Wales. The inequality between the North and South is substantial: as noted by McCann (2020), the UK is now the most regionally unequal industrialised country in terms of GDP per capita, productivity, and disposable income. Closing the North–South wage gap has been a key policy priority in the UK throughout the twenty-first century and is commonly referred to as "Levelling Up." Initiatives to reduce regional disparities have included devolving powers to the Scottish and Welsh governments and relocating government departments—such as the NHS head office—to the North.[25] Reducing regional inequality is also important from a welfare perspective, as improving access to high-quality jobs can increase aggregate output through higher productivity (Bell et al., 2023; Criscuolo et al., 2012; Stansbury et al., 2023).

The government's focus on Levelling Up—alongside the disproportionately large volume of private sector investment in the South—contributed to an elevated share of public sector employment in the North. Figure 12 illustrates this by plotting the share of public employment by postcode region. The public sector share was highest in both Scotland and Wales, peaking at 43% in Swansea. Even in Central London—the seat of the English government—the public share was just 28%, below the 31% national average. Over the same period, wage growth in the North significantly outpaced that in the South. Prior to 2010, wages in the South rose by 4.5 percentage points, while wages in the North increased by 7.4 percentage points, thereby narrowing the regional wage gap.

With these facts in mind, we hypothesise that the higher share of public sector workers in the North led to a greater likelihood of these workers being affected by the pay caps, thereby increasing the North–South wage gap. The left panels of Figure 13 illustrate observed and counterfactual wage growth in each region. As discussed above, wages in the North increased at a relatively faster rate prior to 2010. After 2010, wages declined in both regions, although the decline was larger in magnitude in the South. In the counterfactual scenario in which public sector returns to observable characteristics followed those in the

[25]The NHS head office was moved to Newcastle, and the Office for National Statistics was relocated to Newport, Wales.

private sector, wages would have increased in both regions—but more substantially in the North. The results imply that the pay caps reduced the North–South wage gap by 0.38 percentage points (3.9%).

6 Conclusion

This paper revisits the UK wage distribution during the post-2010 austerity period. We show that overall hourly wage inequality, as measured by the P90–P50 log wage differential, declined by 0.6 percentage points between 2010 and 2018—a reversal of the long-run trend. We propose an institutional explanation for this shift and attribute it to the 2010–2017 public sector pay freeze affecting workers earning above £21,000. Our findings demonstrate that public pay policy has significant distributional consequences—altering the relative position of different groups within the wage distribution and influencing the structure and location of inequality.

The progressive nature of the policy disproportionately compressed wages at the top of the distribution, eliminating the public sector wage premium and driving elevated outflows among higher earners. During the same period, the public sector upskilled more slowly than the private sector—in terms of education, age, and the share of workers in top managerial roles. Additionally, the male share and the proportion of workers based in the South declined more sharply in the public sector, further dampening wage growth relative to the private sector.

In a counterfactual scenario where public sector returns to observable characteristics followed the private sector trend, the P90–P50 wage differential would have increased by 2.9 percentage points rather than declining by 0.6 percentage points. The pay caps also generated unintended distributional spillovers. First, they disproportionately affected women, who constitute the majority of public sector employees. This suppressed female wage growth relative to the pre-policy trajectory and widened the gender pay gap. Second, the caps exacerbated the North–South wage divide, reflecting the higher concentration of public sector employment in the North of England, Scotland, and Wales compared to the

South.

Public sector pay policy has far-reaching economic consequences. Between 2021 and 2023, the UK experienced its largest wave of strikes since 1985, with 5.6 million working days lost—driven largely by disputes over low public sector pay. For example, the average doctor now earns 13% less in real terms than in 2010. These events underscore the broader relevance of our findings: public wage suppression not only shapes inequality but also contributes to labour market unrest. Notably, more than half of OECD countries adopted public sector pay freezes after the financial crisis, suggesting that the dynamics we document extend well beyond the UK.

References

Andrew, A., Bandiera, O., Costa Dias, M., & Landais, C. (2024). Women and Men at Work. *Oxford Open Economics*, *3*, i294–i322.

Andrews, E. (2015, June 17). *Civil Service Staff Numbers Under the Coalition Government, 2010 to 2015*. Institute for Government.

Bell, B., Bloom, N., & Blundell, J. (2022). Income Dynamics in the United Kingdom and the Impact of the Covid-19 Recession. *Quantitative Economics*, *13*(4), 1849–1878.

Bell, B., & Van Reenen, J. (2014). Bankers and Their Bonuses. *The Economic Journal*, *124*(574), 1–21.

Bell, T., Clark, T., Fry, E., Kelly, G., & Thwaites, G. (2023, December 4). *Ending Stagnation: A New Economic Strategy for Britain*.

Blanchflower, D. G., & Bryson, A. (2004). The Union Wage Premium in the US and the UK. *Centre for Economic Performance Working Paper*, *612*(1), 38.

Blundell, R., Crawford, C., & Jin, W. (2014). What Can Wages and Employment Tell Us about the UK's Productivity Puzzle? *The Economic Journal*, *124*(576), 377–407.

Boileau, B., O'Brien, L., & Zaranko, B. (2022). *IFS Green Budget: Public Spending, Pay and Pensions*. Institute for Fiscal Studies.

Bradley, J., Postel-Vinay, F., & Turon, H. (2017). Public Sector Wage Policy and Labor Market Equilibrium: A Structural Model. *Journal of the European Economic Association*, *15*(6), 1214–1257.

Britton, J., & Propper, C. (2016). Teacher pay and school productivity: Exploiting wage regulation. *Journal of Public Economics*, *133*, 75–89.

Burgess, S., Propper, C., Ratto, M., & Tominey, E. (2017). Incentives in the Public Sector: Evidence from a Government Agency. *The Economic Journal*, *127*(605), F117–F141.

Cribb, J., Joyce, R., & Wernham, T. (2023). Twenty-five Years of Income Inequality in Britain: the Role of Wages, Household Earnings and Redistribution. *Fiscal Studies*, *44*(3), 251–274.

Cribb, J., & O'Brien, L. (2024). *Recent trends in public sector pay*. IFS Report.

Criscuolo, C., Martin, R., Overman, H., & Van Reenen, J. (2012). The Causal Effects of an Industrial Policy.

Di Porto, E., Dustmann, C., Incoronato, L., & Gianneto, C. (2023). The effects of salary caps for public managers: evidence from Italy. *Working Paper*.

Dickens, R., & Manning, A. (2004). Has the National Minimum Wage Reduced UK Wage Inequality? *Journal of the Royal Statistical Society Series A: Statistics in Society*, *167*(4), 613–626.

DiNardo, J., Fortin, N. M., & Lemieux, T. (1996). Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach. *Econometrica*, *64*(5), 1001–1044.

Dube, A. (2019). *Impacts of Minimum Wages: Review of the International Evidence*. His Majesty's Treasury.

Dustmann, C., Ludsteck, J., & Schönberg, U. (2009). Revisiting the German Wage Structure. *The Quarterly Journal of Economics*, *124*(2), 843–881.

Emmerson, C. (2017). *Two Parliaments of Pain: the UK Public Finances 2010 to 2017*. Institute for Fiscal Studies.

Facchetti, E. (2024). Police Infrastructure, Police Performance, and Crime: Evidence from Austerity Cuts. *Institute for Fiscal Studies Working Paper*, *24*(16), 71.

Fenizia, A. (2022). Managers and Productivity in the Public Sector. *Econometrica*, *90*(3), 1063–1084.

Fenizia, A., Checchi, D., & Lucifora, C. (2024). Public- and private-sector jobs: a cross-country perspective. *Oxford Economic Papers*, *76*(3), 759–779.

Fetzer, T. (2019). Did Austerity Cause Brexit? *American Economic Review*, *109*(11), 3849–86.

Garibaldi, P., Gomes, P., & Sopraseuth, T. (2021). Public Employment Redux. *Journal of Government and Economics*, *1*, 100003.

Giupponi, G., & Machin, S. (2022). Company Wage Policy in a Low-Wage Labor Market. *Centre for Economic Performance Working Paper*, *1869*(1), 78.

Giupponi, G., & Machin, S. (2024). Labour Market Inequality. *Oxford Open Economics*, *3*, i884–i905.

Gregg, P., Machin, S., & Fernández-Salgado, M. (2014). Real Wages and Unemployment in the Big Squeeze. *The Economic Journal*, *124*(576), 408–432.

Juhn, C., Murphy, K. M., & Pierce, B. (1993). Wage Inequality and the Rise in Returns to Skill. *Journal of Political Economy*, *101*, 410–441.

Lemieux, T. (2006). Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill? *The American Economic Review*, *96*(3), 461–498.

McCann, P. (2020). Perceptions of Regional Inequality and the Geography of Discontent: Insights from the UK. *Regional Studies*, *54*(2), 256–267.

Michael, M., & Christofides, L. N. (2020). The impact of austerity measures on the public - private sector wage gap in Europe. *Labour Economics*, *63*, 101796.

Mulligan, C. B., & Rubinstein, Y. (2008). Selection, Investment, and Women's Relative Wages Over Time. *The Quarterly Journal of Economics*, *123*(3), 1061–1110.

OECD. (2025, June 19). *Government at a Glance 2025*.

ONS. (2019). *What are the Vacancy Trends in the Public Sector?* Office for National Statistics.

Ortiz, I., Cummins, M., Capaldo, J., & Karunanethy, K. (2015). The Decade of Adjustment: A Review of Austerity Trends 2010-2020 in 187 Countries.

Overman, H. G., & Xu, X. (2024). Spatial Disparities Across Labour Markets. *Oxford Open Economics*, *3*, i585–i610.

Postel-Vinay, F. (2015). *Does it Pay to be a Public-Sector Employee?* (No. 156). Institute for the Study of Labor (IZA).

Sepahsalari, A., & Postel-Vinay, F. (2023). Labour Mobility and Earnings in the UK, 1992-2017. *Economic Journal*, (133), 3071–3098.

Singleton, C. (2019). The Public–Private Sector Wage Differential in the UK: Evidence from Longitudinal Employer–Employee Data. *Economics Letters*, *174*, 109–113.

Stansbury, A., Turner, D., & Balls, E. (2023). Tackling the UK's Regional Economic Inequality: Binding Constraints and Avenues for Policy Intervention. *Contemporary Social Science*, 1–39.

Villa, C. (2024). The Effects of Youth Clubs on education and Crime. *Institute for Fiscal Studies Working Paper*, *51*(24), 84.

Figures

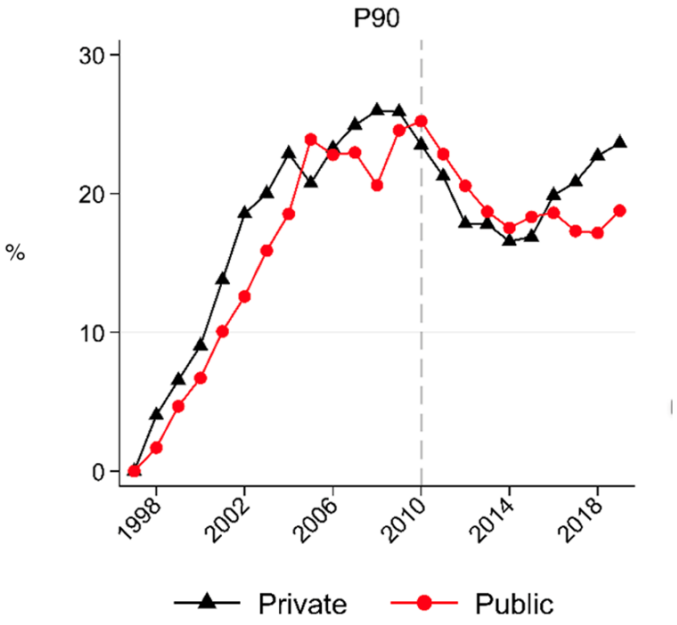

Figure 1: Real P90 Wages Indexed to 1997

Note: Based on ASHE data for public and private sector workers aged 16 and over. The figure plots the evolution of real log hourly wages at the 90th percentile, multiplied by 100 for readability. Wages are indexed to 1997 = 0, so the values represent percentage changes in real wages relative to 1997. All wages are deflated using the Consumer Price Index (CPI), expressed in 2019 prices.

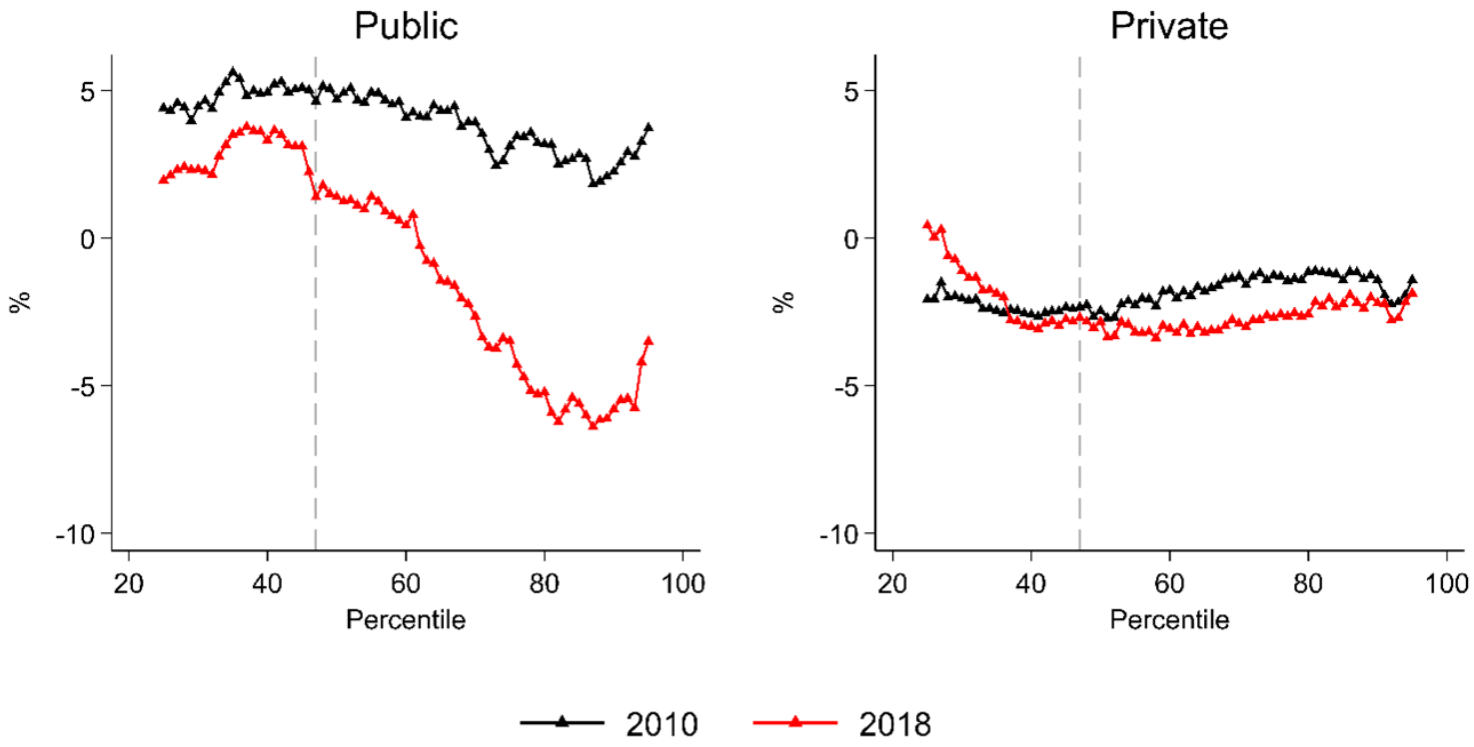

Figure 2: Wage Growth Along the Distribution Relative to 2004

Note: Based on ASHE data for public and private sector workers aged 16 and over. The figure shows real hourly wage growth at equally spaced percentiles, relative to 2004. The vertical dashed line marks the point in the public sector distribution where the pay caps bind. Percentiles below the 25th are excluded to abstract from spillover effects of the living wage; observations above the 95th percentile are omitted.

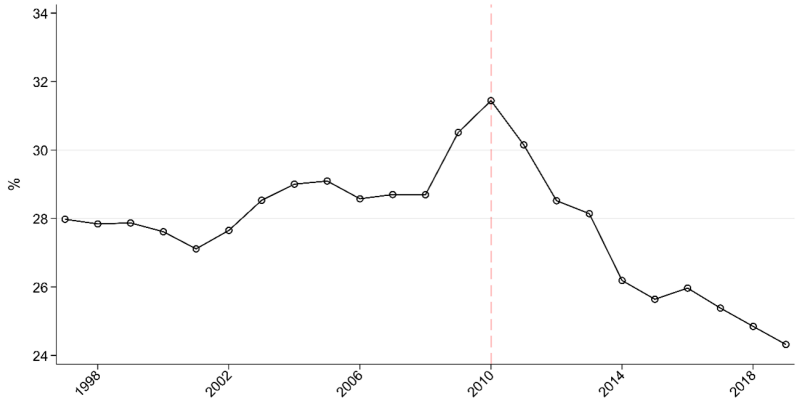

Figure 3: Public Sector Share of Total Employment (1997–2019)

Note: Based on ASHE data for public and private sector workers aged 16 and over. The figure plots the public sector share of total employment between 1997 and 2019.

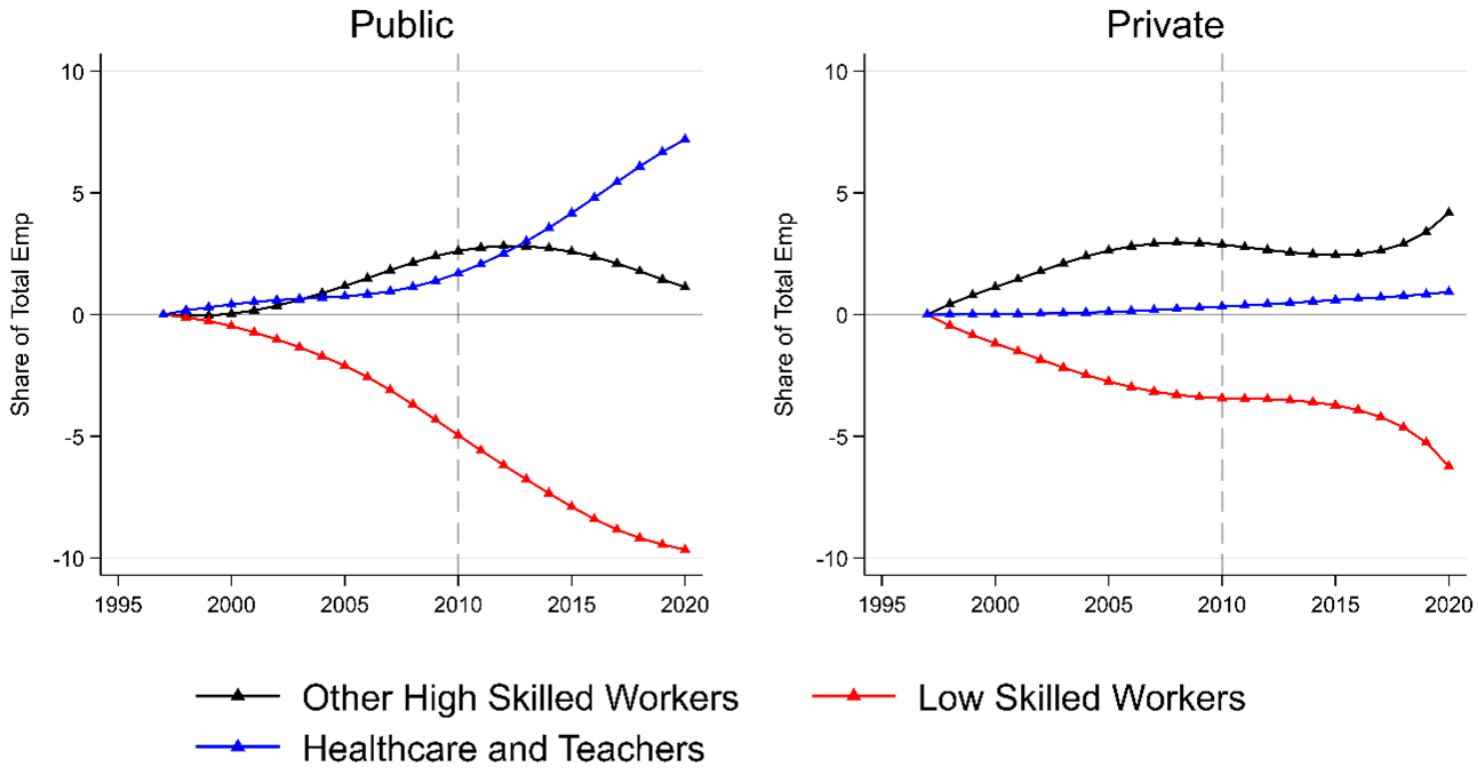

Figure 4: Smoothed Share of Employment by Sector and Occupation

Note: Based on ASHE data for public and private sector workers aged 16 and over. The figure shows the change in the share of total employment over time, normalised to 1997 = 0. The vertical dashed line at 2010 indicates the introduction of the public sector pay caps. High-skilled occupations are defined as those in which the share of degree-educated workers, as a proportion of within-occupation employment, is above the sample mean.

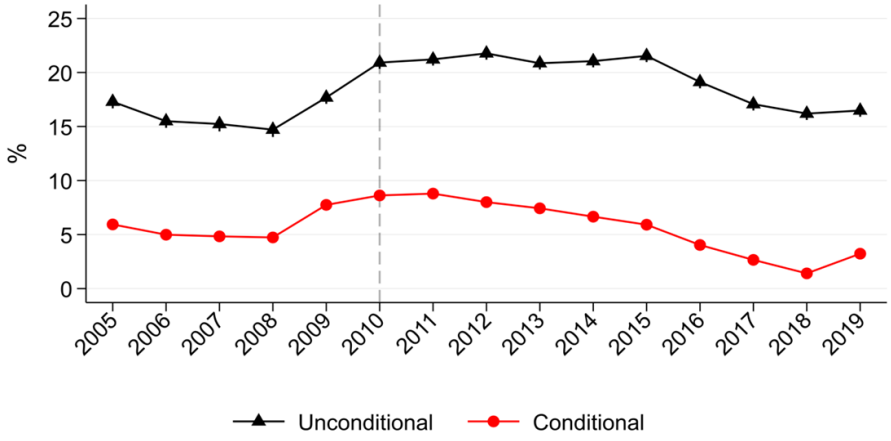

(a) Conditional and Unconditional Premium

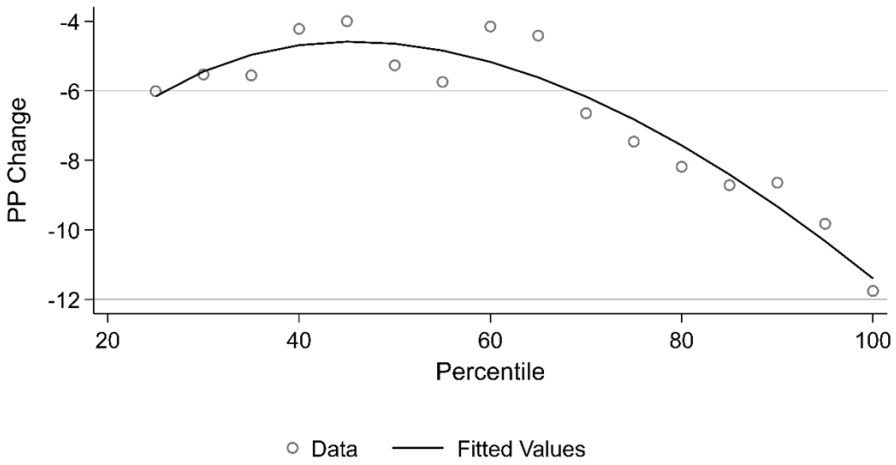

(b) Change in the Conditional Premium Along the Wage Distribution

Figure 5: Public Sector Wage Premium Over Time and Across the Distribution

Note: Based on ASHE data for public and private sector workers aged 16 and over. Panel A shows the unconditional public sector wage premium, calculated as the coefficient on a public sector indicator regressed on log real hourly wages. The conditional premium additionally controls for age (squared), occupation, union status, their interactions, and includes individual and time fixed effects. Vertical lines denote 95% confidence intervals. Panel B shows changes in the conditional premium across 20 wage bins between 2010 and 2018, estimated using a flexible polynomial.

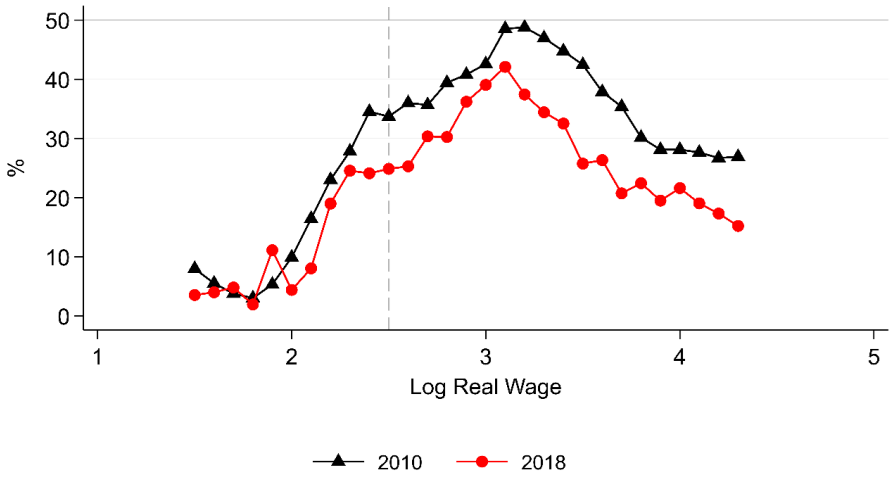

Figure 6: Public Share of Employment Along the Wage Distribution

Note: Based on ASHE data for public and private sector workers aged 16 and over. The figure uses the unconditional wage distribution in 2010 and 2018 to compute the public sector share of employment across the wage distribution in 0.1 log-point intervals. Wages are measured hourly and deflated to 2019 prices. The vertical line marks the pay cap threshold.

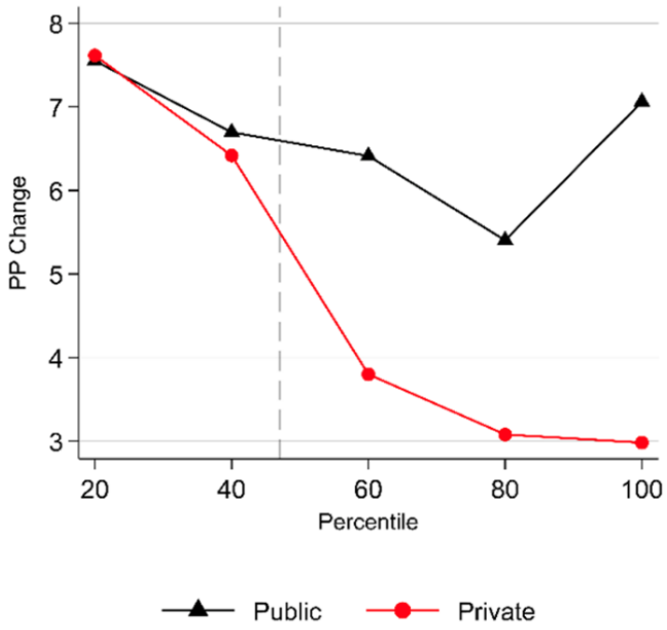

Figure 7: Change in Outflow Along the Wage Distribution, 2010–2018

Note: Based on ASHE data for public and private sector workers aged 16 and over. The figure plots the percentage point change in annual outflow rates across the wage distribution, separately for the public and private sectors. The initial period pools 2008–2009, and the later period pools 2017–2018. Wages are deflated to 2019 prices. Vertical lines indicate the binding threshold of the public sector pay caps.

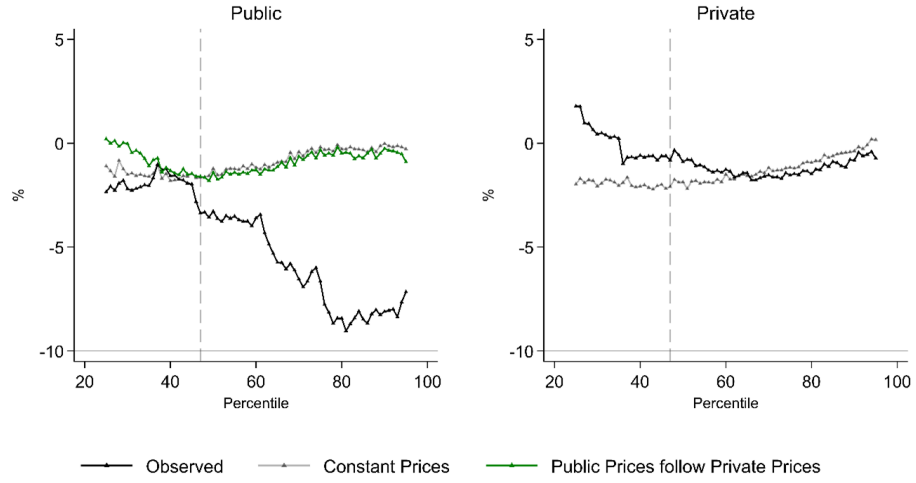

(a) Price and Relative Price Effects

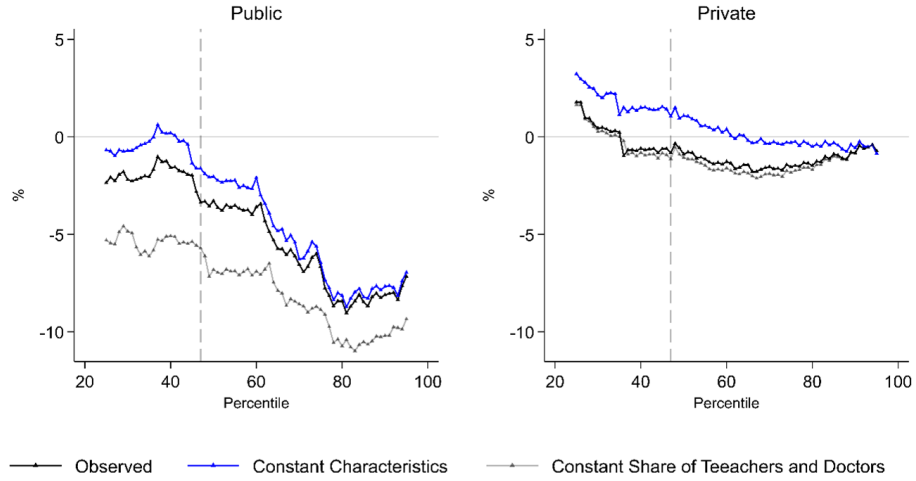

(b) Characteristic Effects

Figure 8: Observed and Counterfactual Wage Growth Along the Distribution (2010-18)

Note: Based on ASHE sample for public and private sector workers aged 16 and over. Both panels show observed and counterfactual real wage growth between 2010 and 2018. Panel A compares the observed growth (black line) with two counterfactuals: (i) wages if returns to observable characteristics had remained fixed at 2010 levels and characteristics evolved as observed in 2018 (grey line), and (ii) wages if returns to observable characteristics in the public sector had changed in line with those in the private sector (green line). Panel B compares observed wage growth (black line) with counterfactuals holding constant: (i) the distribution of characteristics at 2010 levels (blue line), and (ii) the 2010 share of teaching and medical professionals, while allowing other characteristics to follow their 2018 distribution (grey line). Counterfactuals are computed using a DiNardo et al. (1996) decomposition, where wages are modelled using a linear specification including sex, age (squared), union status, hours worked, and region and occupation fixed effects. The vertical dashed line denotes the bite point of the pay caps.

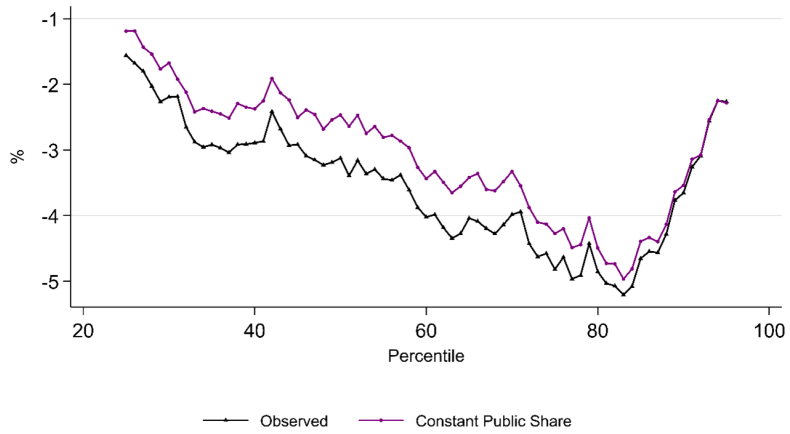

Figure 9: Observed and Counterfactual Wage Growth Along the Distribution Assuming a Constant Public Sector Share (2010–2018)

Note: Based on ASHE sample for public and private sector workers aged 16 and over. The figure shows observed and counterfactual wage growth between 2010 and 2018. The black line represents observed wage growth; the purple line represents the wage growth that would have occurred if the public sector share had remained fixed at 2010 levels. Counterfactuals are calculated via a DiNardo et al. (1996) decomposition, where wages are generated using a linear specification including sex, age (squared), union status, hours worked, and region and occupation fixed effects.

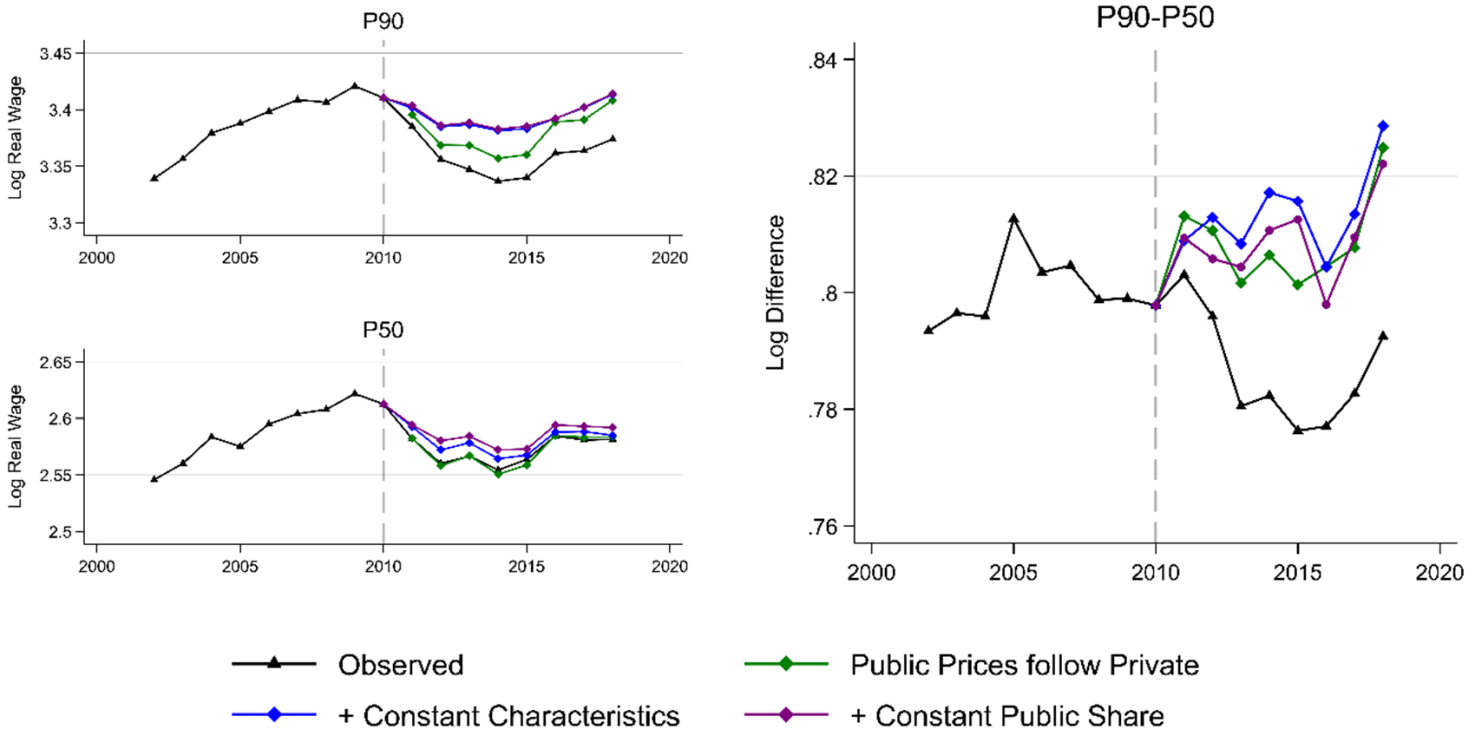

Figure 10: Counterfactual P90–P50 Inequality for the Whole Economy

Note: Based on ASHE sample for public and private sector workers aged 16 and over. The figure displays observed and counterfactual inequality, measured by the difference between the 90th and 50th percentiles of the wage distribution. The difference between the black and green lines represents the wage distribution that would have prevailed had public sector returns to observable characteristics increased in line with those in the private sector. The difference between the green and blue lines captures the effect of changes in public sector characteristics. The difference between the blue and purple lines reflects the additional impact of holding the public sector employment share constant. Counterfactuals are calculated using the decomposition method of DiNardo et al. (1996), where wages are generated using a linear specification that includes sex, age (squared), union status, hours worked, and region and occupation fixed effects.

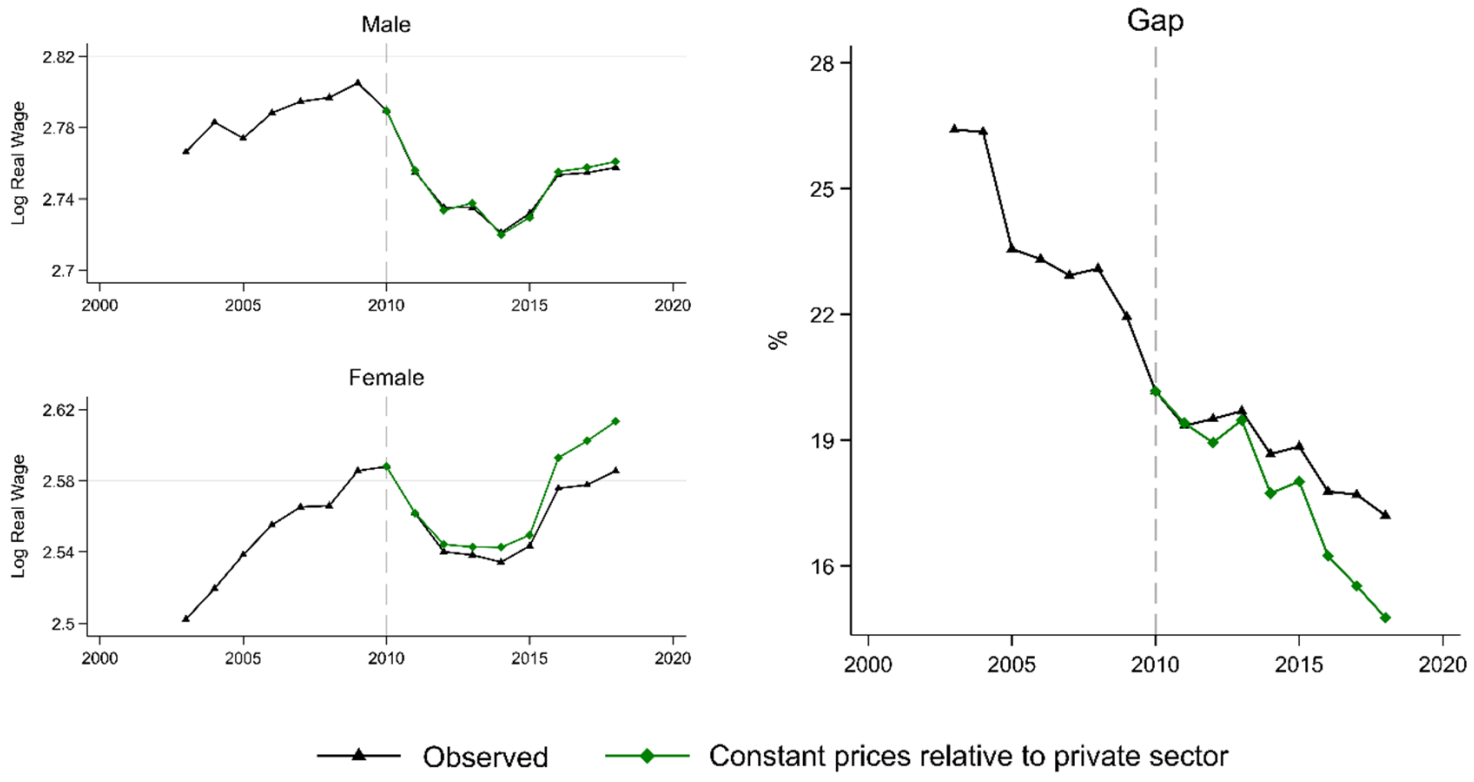

Figure 11: Counterfactual Gender Pay Gap

Notes: Based on ASHE sample of public and private sector workers aged 16 and over. The figure displays observed and counterfactual inequality, measured by the difference between the mean male and mean female wage. The difference between the black and green lines represents the wage distribution that would have prevailed had public sector returns to observable characteristics increased in line with those in the private sector. Counterfactuals are calculated using the decomposition method of DiNardo et al. (1996), where wages are generated using a linear specification that includes sex, age (squared), union status, hours worked, and region and occupation fixed effects.

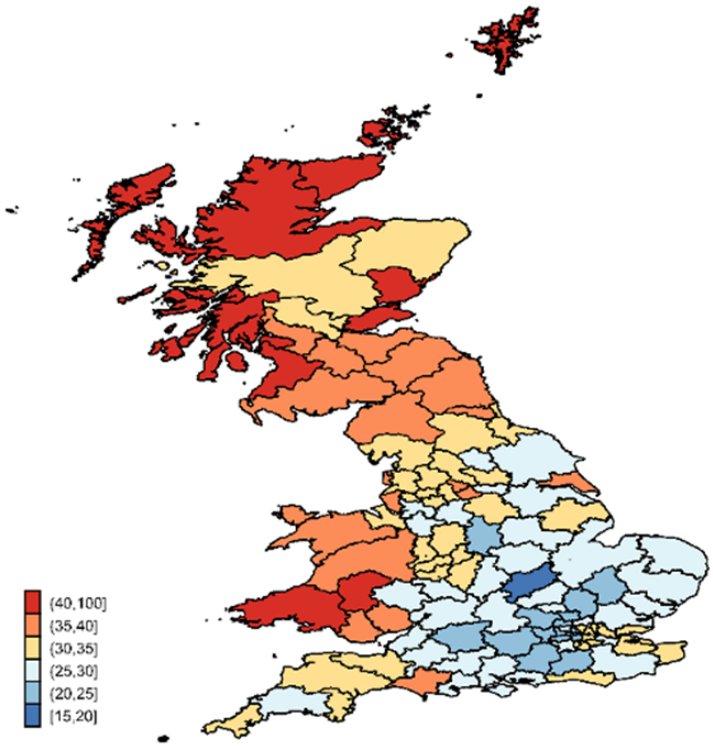

Figure 12: Public Sector Share by Region (2010)

Note: Based on ASHE sample of public sector workers aged 16 and over. The figure shows the public sector share of employment at the postcode region level. The public share refers to the percentage of employed workers in the public sector in 2010.

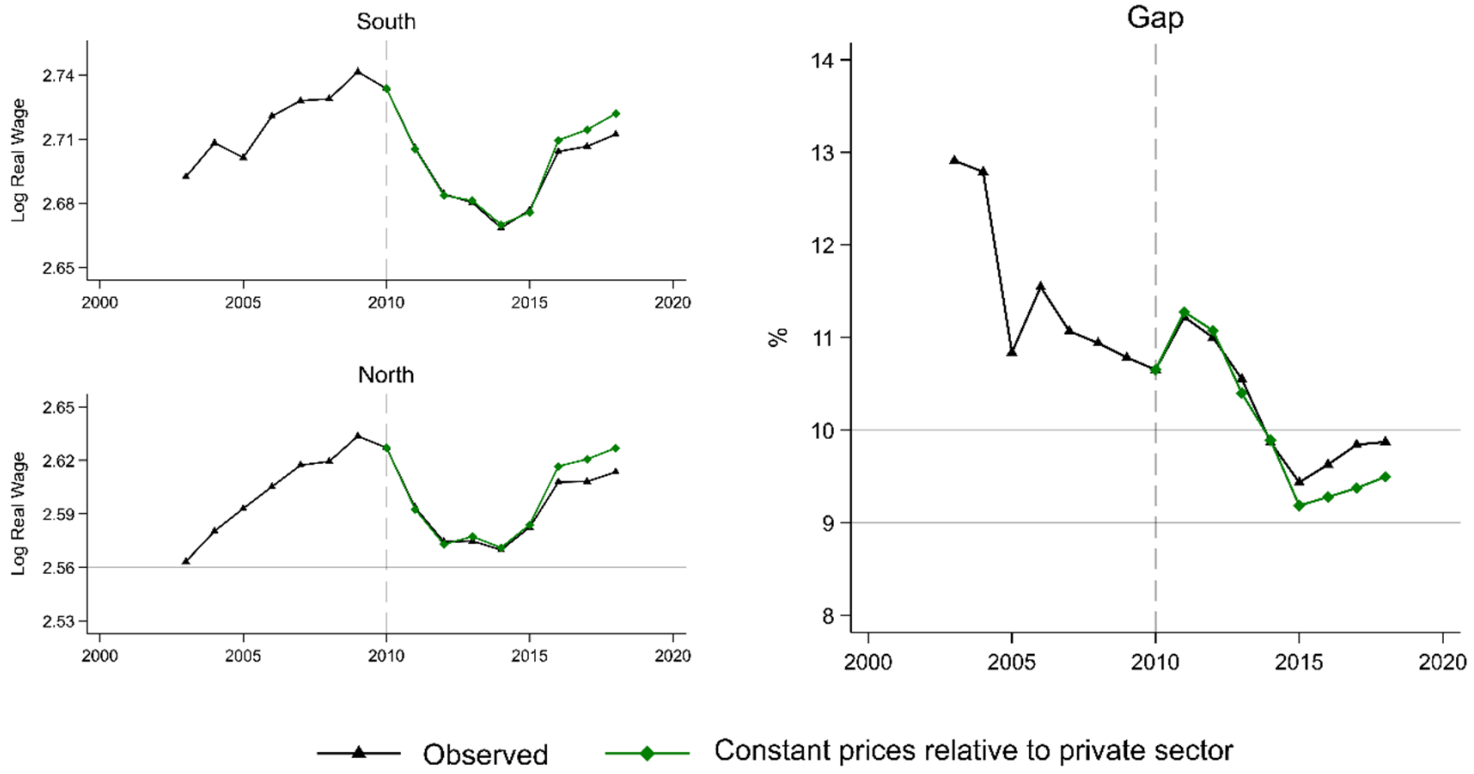

Figure 13: Counterfactual North–South Wage Gap

Notes: Based on ASHE sample of public and private sector workers aged 16 and over. The figure shows observed and counterfactual North–South wage inequality, measured as the difference between the mean wages in the North and South. The black line represents observed data; the green line shows how earnings and the regional wage gap would have evolved if public sector returns to observable characteristics had changed in line with those in the private sector. Counterfactuals are calculated using the decomposition method of DiNardo et al. (1996), where wages are generated using a linear specification including sex, age (squared), union status, hours worked, and region and occupation fixed effects. The North is defined at the government office region level and includes the North East, North West, Yorkshire and the Humber, Wales, Scotland, and Northern Ireland.

Tables

Table 1: Public and Private Summary Statistics

	Public			Private		
	(1)	(2)	(3)	(4)	(5)	(6)
	2010	2018	% Chg	2010	2018	% Chg
Share of Employment	31.6	24.0	-24.7***	68.8	76.1	10.3***
Mean Hourly Wage	2.79	2.74	-1.9***	2.58	2.57	-0.4***
P90	3.40	3.31	-2.7***	3.32	3.28	-1.3***
P10	2.21	2.20	-0.5***	1.98	2.08	5.1***
Degree^{+}	57.0	64.9	14.0***	32.9	41.9	27.3***
Male	34.1	30.1	-11.9***	57.1	55.3	-3.1***
Age $\leq$ 25	6.6	6.6	-0.4	16.8	16.2	-3.9***
(26,55]	78.1	75.0	-4.0***	69.8	67.9	-2.7***
55+	15.3	18.4	20.3***	13.4	15.9	18.9***
North	43.8	45.2	3.2***	35.3	35.6	0.7
Weekly Hours	31.8	31.7	-0.4	34.6	34.1	-1.4***
Observations	46883	36167	83050	102001	113633	215634

Note: Author's calculations based on the ASHE sample of workers aged 16 and over. The table reports selected summary statistics for individuals employed in the public and private sectors. Hourly wages are deflated to 2019 prices. Workers in the North are defined using government office regions and include the North East, North West, Yorkshire and the Humber, Wales, Scotland, and Northern Ireland. $^{+}$ Degree share is calculated using the UKLS; a degree-educated worker is defined as anyone with education beyond A-levels.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Transition Matrices

		t+1					
		Public	Private	Self Emp	Unemp	Retired	Inactive
	Panel A: 2004-2009						
t	Public	88.0	5.0	0.9	1.0	1.9	3.3
	Private	2.6	87.0	2.9	2.4	1.6	3.5
	Self Employment	1.5	11.7	81.4	1.5	1.7	2.1
	Unemployment	6.8	27.5	5.5	32.9	3.5	23.9
	Retired	0.3	0.8	0.2	0.2	95.6	2.9
	Inactive	5.3	12.4	1.5	5.0	5.3	70.5
	Panel B: 2011-2018						
t	Public	79.9	12.1	1.1	1.2	2.9	2.8
	Private	5.8	83.0	3.5	2.4	1.9	3.5
	Self Employment	1.5	12.7	78.7	1.3	3.5	2.3
	Unemployment	4.4	21.9	4.1	42.9	3.5	23.3
	Retired	0.3	0.6	0.5	0.2	96.8	1.6
	Inactive	4.2	12.6	1.7	8.4	4.3	68.9
	Panel C: Difference						
t	Public	-8.1	7.1	0.2	0.2	1.0	-0.5
	Private	3.2	-4.0	0.6	0.0	0.2	0.0
	Self Employment	0.0	1.0	-2.7	-0.2	1.8	0.2
	Unemployment	-2.4	-5.6	-1.4	10	0.0	-0.6
	Retired	0.0	-0.2	0.3	0.0	1.2	-1.3
	Inactive	-1.1	0.2	0.2	3.4	-1.0	-1.6

Note: Labour market transitions based on UKLS data. Individuals are classified into six labour market states: public sector employment, private sector employment, self-employment, unemployment, retirement, and inactivity. Panel A reports average annual transition probabilities between states from 2004 to 2009. Panel B reports the same for 2011 to 2018. Each row shows the probability of being in a given state in year $t+1$, conditional on the state occupied in year t; thus, row probabilities sum to 100. Panel C presents the percentage point change in transition probabilities between the two periods. The year 2010 is excluded due to the transition from the British Household Panel Survey (BHPS) to Understanding Society (USoC), which involved a change in sampling structure. We reweight the sample to produce a balanced cross-section and mitigate composition effects due to ageing.

Appendix Figures & Tables

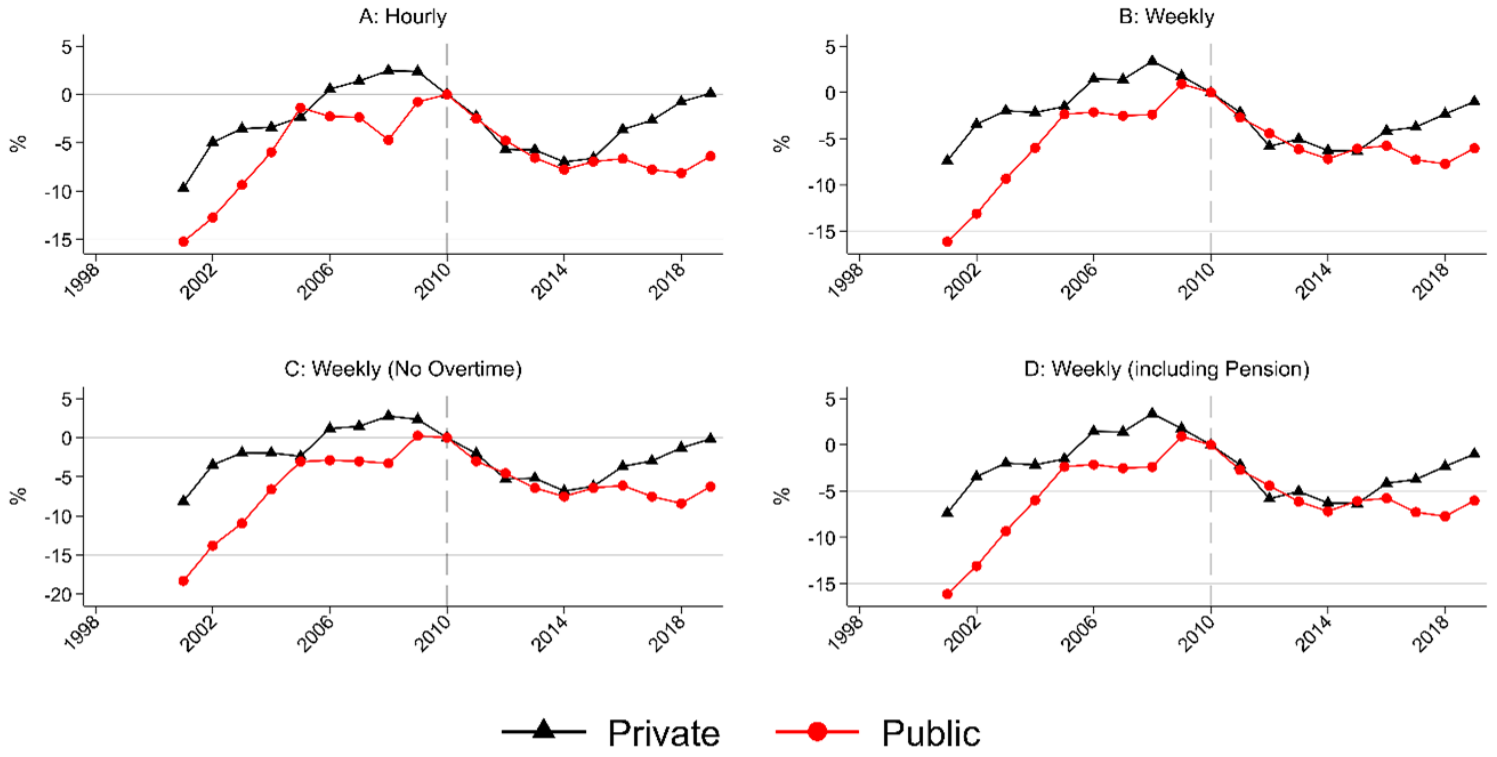

Figure A1: P90 Real Earnings Robustness Checks Normalised to 2010

Note: Based on ASHE data for public and private sector workers aged 16 and over. The figure displays alternative measures of log real wages, all normalised to 2010 = 0. Panel A shows the baseline hourly wage measure used throughout the analysis. Panel B reports weekly wages to capture intensive margin effects. Panel C excludes overtime pay, and Panel D includes employer pension contributions.

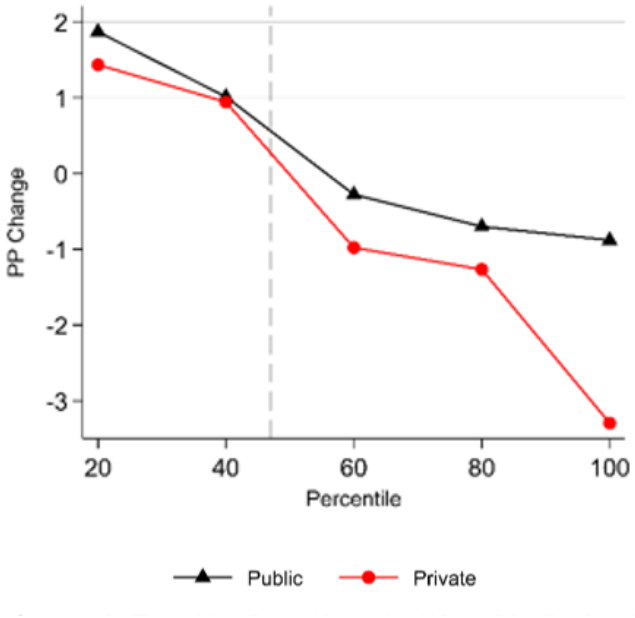

Figure A2: Change in Inflow Along the Wage Distribution, 2010–2018

Note: Based on ASHE data for public and private sector workers aged 16 and over. The figure plots the percentage point change in annual inflow rates across the wage distribution, separately for the public and private sectors. The initial period pools 2008–2009, and the later period pools 2017–2018. Wages are deflated to 2019 prices. Vertical lines indicate the binding threshold of the public sector pay caps.

Table A1: Summary of Decomposition Effects

(1) Effect	(2) Sub-Effect	(3) Definition	(4) Factor Held Constant	(5) Sample
Characteristic	-	Change in wage distribution due to changes in the distribution of observable worker characteristics	Prices (β)	Public only
	Occupational Composition	Change in wages due to shifts in the share of teaching/medical professions	Prices (β), other covariates	Public only
	Other Composition	Change in wages due to shifts in the distribution of other observable characteristics	Prices (β), occupation	Public only
Price	-	Change in wage distribution due to evolving returns to observable characteristics between 2010 and 2018	Composition (x)	Public only
	Relative Price	Change in the wage distribution due to changing returns to observable characteristics in the public sector relative to the private sector	Composition (x)	Public only
	Residual Price	Change in the wage distribution due to public sector price changes not explained by the private sector counterfactual	Composition (x)	Public only
Share	-	Change in overall wage distribution due to falling public sector share of employment	Prices (β), Composition (x)	Public and private sector

Note: Table of counterfactual effects identified throughout Section 4. Column (1) defines the effect, and column (2) identifies any sub-effects within this group. Column (3) explains the effect; column (4) outlines the component of the distribution held fixed. Column (5) states the data used to identify the effect.