# The UK Graduate Premium in the Last Quarter Century: Patterns and Driving Forces*

Richard Blundell, David A. Green, Wenchao Jin

October 2026

**Abstract**

The UK graduate wage premium has been falling since 2010, a pattern widely read as evidence that the expansion of higher education has outstripped demand for graduate skills. We show that this reading is hard to sustain. The premium is flat until 2010, and the subsequent decline is predominantly a non-London phenomenon. Relative supply effects cannot account for the post-2010 declines outside London: the regions where the graduate share grew fastest are not those where the premium fell most. The decline is driven by a relative increase in the real wage of workers whose highest qualification is a good secondary-school credential (HS graduates) outside versus inside London rather than by a relative decline in the university graduate wage. We show that the relative increase in the HS wage is consistent with spillover effects from the substantial increases in the National Minimum Wage after 2015. This has the potential to shift debates about the implications of the post-2010 decline in the graduate wage premium.

**JEL codes:** J31, J38

## 1 Introduction

The decline in the UK graduate wage premium has raised considerable concern in recent years. Framed in a standard demand and supply discussion, researchers and analysts have

*Blundell (UCL, IFS, TSE, r.blundell@ucl.ac.uk); Green (UBC, IFS, david.green@ubc.ca); Jin (Sussex, IFS, Wenchao.Jin@sussex.ac.uk). We would like to thank Attila Lindner and Xiaowei Xu for helpful comments. We acknowledge support from ESRC (CPP at IFS, ES/M010147/1). Analysis was carried out in the ONS Secure Research Service. Statistical data © Crown Copyright; the ONS bears no responsibility for interpretation or analysis.

argued that the very substantial increase in higher education in the UK since the mid-1990s has resulted in the supply of university level skills outstripping demand (Stansbury et al., 2023; Burn-Murdoch, 2026). As a direct measure of this, Xu (2023) shows that about a third of BA graduates in London work at jobs that do not require a BA, and the percentage is higher outside London. The appropriate policy response has variously been described as to either cut back on the supply or find a way to improve the demand for skills. In either case, the declining trend in the graduate earnings premium is viewed as a bellwether for deep problems with human capital in the UK. In this paper, we take a close look at movements in the wage premium paid for university degrees in the UK. Our belief is that a careful depiction of the source and direction of movements in the graduate premium can help in understanding what the premium is actually telling us, potentially redirecting policy discussions.

The first step in our investigation is to establish the nature of the trend in the graduate wage premium, which we will define as the ratio of the median wage paid for workers with a bachelor or higher degree from a university ("BA's") to the median wage for workers with a good secondary level education: "high school graduates" ("HS's").[1] Perhaps surprisingly for such a simple construct, different analysts have presented graduate premium series with quite different trends. We show that the biggest factor driving differences in the trend is whether the lower education group to which university graduates are compared includes the bottom education group, those without good GCSEs. Including that lowest group results in a declining trend for the wage premium reaching back to the 1990s. Including only the "high school graduates" produces a flat trend in the premium up to 2010 in all regions, continued flatness for London after 2010, and a decline outside London after 2010; the UK-wide version of this series reproduces Blundell et al. (2022).

Whether we want to focus on the series that includes everyone with a below university education as the comparator group or restrict attention to the high school graduates depends, in part, on the question we are seeking to investigate when we use the premium series. An interest in understanding how education relates to movements in inequality might point to working with the first definition since we may be interested in how the growing university educated population is faring relative to everyone else. But movements in education differentials are also useful because common models of technological change are built on the assumption that new technologies are differentially complementary with different skill levels. Given that, movements in wage-skill ratios can shed light on the nature and extent of technological change. In this case, we would want to use more homogeneous skill groups whose wages could be seen as the prices for specific factors of production. Based on this and a concern that wages for those without good GCSEs may reflect minimum wage movements

[1] A good GCSE level means at least a grade C (grade 4 in the current 9–1 system) in GCSE qualifications, without going on to complete a degree.

to a substantial degree, we argue for focusing attention on the wage ratio between university graduates and what we call high school graduates. It is that ratio that follows the flat then declining pattern outside of London.

This characterisation of the wage premium patterns refocuses attention on the question, why did the wage premia inside and outside London part company after 2010? One explanation that has been advanced is that, with universities expanding across the UK, supply increased in all regions but demand only kept up with supply in London. In other regions, university graduates could only find jobs that did not require a university degree (Xu, 2023). If this lack of ability to find graduate jobs is a proximal cause, such a finding would push us to ask why the regions outside of London could not keep up in creating demand for high skilled workers which, in turn, links to investigations of the relative differences in productivity growth inside and outside London.

We investigate the supply and demand explanation for the differential movement in the premium inside and outside London in two ways. First, we plot the BA/HS wage premium against the BA share of employment across region by three-year-period cells, after residualising both on a complete set of region and period fixed effects so that each is measured relative to the UK-wide trend. That exercise does not support the relative demand and supply hypothesis: regions with higher growth in the proportion of workers with a BA over the period from 1997 to 2025 did not show noticeably different movements in the BA/HS wage premium. We also follow Xu (2023) in investigating the under-placement of university graduates inside and outside London. We find that there is more occupational downgrading of graduates outside London versus in London since 2001 but that trend had virtually stopped by 2010. Fitting with the lack of relative downgrading after 2010, once wages are age-adjusted, graduate wages outside London show no significant decline relative to London after 2010. Thus, under-placement can account for little of the relative decline in the university premium outside relative to inside London. Instead, the fall in the wage premium outside London after 2010 comes from a somewhat surprising source: a rise in HS wages outside London at a time when HS wages in London are falling along with BA wages. This serves to further re-focus the relevant question to, why did HS wages behave so differently outside versus in London after 2010?

We argue that the most likely source of the different movements in HS wages outside versus in London is the differential impact of minimum wages. Beginning in 2016, successive UK governments raised the National Minimum Wage (NMW) substantially over a number of years. In nominal terms, the median private-sector HS wage in 9 of the 12 UK regions was at least £3 above the NMW in 2010. In 2016, instead, 7 of 12 regions had median HS wages that were only between £2 and £3 above the NMW, and 2 regions had median HS

wage £1-2 above the NMW. That would place the median HS wages in those regions within a range identified as being affected by spillovers from minimum wages in the small literature examining those spillovers (Cengiz et al., 2019; Gopalan et al., 2021; Giupponi et al., 2024; Brochu et al., 2025). Using both regional and occupation x region variation and two different datasets, we show that median HS wages move with a minimum wage "bite" variable that is constructed as the NMW value in a year divided by the median HS wage in a region in a base year (2010). Under the assumption that a region's 2010 base median does not otherwise predict its later wage growth, the bite accounts statistically for most of the relative rise in HS wages outside London. We conclude that spillovers from the minimum wage are the leading candidate explanation for the rise in HS wages outside London, and hence for the decline in the BA/HS wage premium there, although we cannot completely rule out other explanations. In London, where the median HS wage is always at least £5 above the NMW in 2025 prices, no such spillovers would be expected.

These results have the potential to alter conclusions about what movements in the BA/HS premium are telling us. Rather than telling us about demand for skills being unable to keep up with supply, the decline in the premium after 2010 most plausibly reflects the impact of the minimum wage intervention in the labour market. Taken together, the results presented here raise a specific set of questions that are different from the focus of the recent debate. In particular, why did the wage premium not change between 1997 and 2010 for all of the UK and continue not to change after 2010 in London in the face of a very substantial increase in the relative supply of BA's? What would have happened to HS wages outside of London in the absence of minimum wage effects (and would that have implied that the BA/HS wage premium would have stayed flat there as well)? And, what are the impacts of a minimum wage increase that is large enough to alter median HS wages on the adoption of new technologies?

In our earlier work (Blundell et al., 2022), we provide an answer to the question of how a substantial increase in the quantity of skills could have left the skill premium unchanged and evidence supporting that answer. In essence, firms, in the face of a substantial increase in the supply of BA's, switch toward an existing technology that is intensive in its use of BA's. Thus, the education increase effectively creates its own demand. But further questions about what such a process means for innovation and productivity, and what minimum wages that cut deep into the wage distribution imply for both the process and its effects, remain to be investigated. We believe that the wage patterns we present here indicate that this is a potentially fruitful direction for future research.

# 2 Data

## 2.1 The QLFS and ASHE

We draw on two complementary sources of UK earnings data. Our main source is the individual-level Quarterly Labour Force Survey (QLFS), pooling the quarterly QLFS files from 1992Q2 through 2025Q4.[2] Our hourly wage measure is real gross weekly pay in the main job divided by total *actual* hours worked in the main job, expressed in 2025 prices, deflated by the GDP deflator. Because the earnings and hours questions are asked only in waves 1 and 5 of the five-wave rotating panel, the sample of hourly-wage observations is much smaller than the full LFS. In many parts of our investigation, we delineate wage and employment movements by region. We use the official 12 first-level statistical regions of the UK: London, South East, South West, East of England, West Midlands, East Midlands, Yorkshire and the Humber, North West, North East, Scotland, Wales and Northern Ireland.

The LFS wage data carry several caveats that we introduce here and return to below when we assess whether the post-2010 wage patterns are genuine. First, response rates and achieved sample sizes have fallen sharply since the pandemic, raising questions about the reliability of the most recent wage data (Office for National Statistics, 2026c).[3] Linking the LFS to HMRC PAYE Real Time Information records, the ONS estimates that the weighted LFS under-counted PAYE employees by an average of 217,000 per quarter in 2019 and 365,000 in 2022–2024, over-counted them by 332,000 during the pandemic (2020–2021), and was within 10,000 in 2025 (Office for National Statistics, 2026b). The ONS also finds that the LFS overstates hours worked, and the bias grew as response rates fell (Office for National Statistics, 2026a). Importantly for our work, the largest distortions in hours appear to be for those without good GCSEs. Second, the derived hourly wage contains substantial measurement error in the lower tail of the distribution: around 15% of employees record a derived hourly wage below the National Minimum Wage (NMW) in 2019, rising to roughly 17% by 2025 (the share was far smaller in 2015), and the LFS 10th percentile sits below the prevailing NMW in these years whereas this is not the case in the arguably more reliable Annual Survey of Hours and Earnings (ASHE) data that we describe momentarily.

[2]We use the End User Licence (EUL) version distributed by the UK Data Service (Office for National Statistics, 2025b). The QLFS is owned and deposited by the Office for National Statistics (ONS), is Crown Copyright, and is accessed under the UK Data Service (UKDS) End User Licence Agreement. Because our analysis pools a large number of quarters, we cite the QLFS series as a whole; the study numbers (SN) and DOIs for the individual quarterly files are available from the UK Data Service.

[3]Barbas et al. (2026) show that these sampling problems are severe enough to distort headline labour-market aggregates: since 2024 the LFS-based count of employees has diverged sharply from the administrative Real Time Information series, and they argue that administrative data now provide the more reliable picture of recent developments.

Our second data source is the ASHE, published by the ONS (Office for National Statistics, 2025a). Unlike the household-based LFS, ASHE is an administrative, employer-reported survey: it is an approximately 1% sample of employee jobs drawn from HMRC PAYE records, with earnings and hours reported by employers rather than by respondents. It is therefore presumably more accurate on pay than a household survey and is the ONS's authoritative source on earnings. We downloaded the ASHE "Hourly Pay" tables from the ONS at the four-digit SOC occupation by region level for 2002–2025, yielding cells defined by occupation × region × year.[4] The SOC2010 grid that is used in the data from 2011 to 2020 spans 369 four-digit occupations and the 11 Great Britain regions (the nine English regions plus Wales and Scotland; Northern Ireland is not covered), giving about 4,059 potential cells per year over 2012–2020, and roughly 4,500 cells per year starting in 2021 when ASHE switched to the SOC2020. The ONS suppresses estimates for cells whose sampling error is too large, however, so a median wage is published for only 2,900–3,000 cells per year.[5] A key limitation for our purposes is that ASHE contains no information on education or qualifications, so it cannot be used to measure the graduate wage premium directly.

The two datasets therefore play distinct roles. The LFS is our main workhorse: because it records educational qualifications, we use it to measure the university graduate versus high-school wage premium and how it has evolved across regions. ASHE, which lacks education, is used to validate the insights from the LFS—in particular to check LFS wage levels and the behaviour of the lower tail of the wage distribution, and to corroborate the patterns we find at the occupation level. Reassuringly, the overall median wage in the LFS tracks the median in the ASHE data very closely, implying that the LFS data problems mentioned earlier are concentrated in the bottom of the distribution rather than around the medians on which our education-premium series rely. We provide a detailed comparison of wage patterns for HS workers in the two datasets in Section 7.

## 2.2 Defining Education Categories

The UK education system consists of a series of levels. Historically, students could legally leave school at age 16 and take the General Certificate of Secondary Education (or GCSE) exam at that age. Obtaining a grade of C or higher on the exam is generally seen as the

[4]Occupations are coded to the SOC2000 over 2002–2010, to the SOC2010 over 2011–2020 and to the SOC2020 from 2021, so occupation codes are comparable only within each of those three blocks.

[5]Roughly a quarter to a third of cells lack a published median, with the lower percentiles being the most likely to be suppressed. Employment counts are suppressed much more often than medians: of the 36,531 potential cells over 2012–2020, 26,388 carry a published median but only 11,047 carry a published employment count. Our occupation-level regressions need both, plus a published 2012 base median, which is why they run on about ten thousand cells rather than twenty-six thousand.

equivalent of obtaining a good secondary education credential. At roughly age 18, students interested in going to university take a second set of exams - the A levels. An alternative definition of attaining a secondary education is having obtained A levels. Under the UNESCO's International Standard Classification of Education (ISCED 2011), both GCSEs (with a C or higher grade) and A-levels in the UK are classified as level 3 – "upper secondary education" – which puts them in the same level as a high school diploma in the US or Canada. We view the GCSE C grade as the more relevant cut-off because it includes people who obtain a secondary education without necessarily a desire to go to university. The high school graduation designation in other countries does not require students to also be interested in going on to tertiary education.[6] Thinking in terms of production functions, we see a grouping of workers with a good secondary education, not divided up by whether they might have an interest in going to university, as a relevant skill group since university intentions would likely not affect the tasks they are assigned.

Beyond these levels of education, there are various diplomas and degrees and also years of education without obtaining a bachelor's degree or equivalent. We treat those levels of education as part of the high school graduate (HS) group. That is, we define a HS grad as having at least a C grade on the GCSE plus (possibly) A-level exams and other non-tertiary education but not a university degree at the bachelor's level or higher. We group together university degree holders with a bachelor's degree or more advanced, post-graduate degrees. For brevity, we will refer to this group as BA's, though they include people with bachelor level degrees other than BA's as well as post-graduate degrees. We present evidence on movements in the BA relative to the HS wage using different definitions of who is a high school graduate but always with the same definition of BA's.

The age composition of the workforce plays an important role in uncovering education related wage differentials. The combination of the ageing of the baby boom and the educational expansion after 1990 in the UK means that the experience level of workers with different education levels can differ over time. Thus, the BA group grew rapidly after 1990 and, so, had an average level of experience that declined over time. With older generations being more concentrated in the HS and lower education levels, the average experience level for those groups rose over time as the baby boom got older but then started to decline as they retired. The result is that simple comparisons of mean or median wages for BA and HS workers confound education with experience effects. To address this, we compute median real hourly wages, using the LFS earnings weights, within cells defined by year, education group and five-year age band (20–24 to 55–59). We then regress the log BA/HS ratio of these

[6] In fact, using the GCSE C grade cut-off among 25-29 year olds in 2012, the proportion we would classify as high school graduates is 53%, which compares favourably to the 56% of US young people in the same age group and year who are classified as high school graduates (Blundell et al., 2022).

cell medians, unweighted, on age-band and year dummies. The year coefficients, normalised to zero in a reference year, are our age-adjusted graduate premium. Age-adjusted wages for a single education group are obtained in the same way from the log of that group's median. For the region-level regressions (Tables 3 and 4), this is done separately for each region and education group, with cells weighted by employment. Appendix A shows that age-adjusting instead with an individual-level Mincerian log-wage regression, as in Stansbury et al. (2023), gives the same patterns.

# 3 Wage Ratio Trends

## 3.1 UK Wide

We begin by showing the log ratio of BA to HS median wages for the UK as a whole for 1993 to 2025 using LFS data in figure 1. This is a replication of figure 2 in Blundell et al. (2022) (hereafter BGJ) up to 2016 and an extension through 2025. As in that paper, one can see that the wage ratio is flat from 1993 through to about 2010. It then dips and rebounds over 2011–2014, leaving the 2014 value at essentially the same level as in 1993, before declining relatively continuously through 2025, at which point it is approximately 9% below its 1993–2010 level. As the next subsection shows, the decline is predominantly a non-London phenomenon that begins around 2010 but is partly masked in the UK aggregate by the flat London series until the middle of the decade. These patterns were accompanied by ongoing increases in the percentage of the population in our BA educational category, shown in panel (b) of Figure 1. That percentage among 20–59 year olds was 11% in 1993, 26% in 2010, and 46% in 2025, while the share with below-GCSE C qualifications fell correspondingly from about a third to under a tenth.

In BGJ, we focused on the somewhat puzzling combination of the substantial increase in the proportion of workers with a BA with the lack of movement in the BA to HS wage ratio over the period up to about 2014. We presented evidence that this fits with an endogenous technological choice model in which the increased supply of BA workers induced firms to switch to an existing, newer technology that makes intensive use of BA workers. During the course of the switch-over, the added supply draws out enough demand to leave the wage ratio unchanged. The model is essentially a variant of a trade model with different technologies taking the place of different goods. Those models yield a factor price invariance result as long as the economy is in the "cone of diversification" - i.e., the switch-over period. Once the whole economy has adopted the new technology, however, the economy reverts to a simple neoclassical form and further increases in BA supply would lead to a reduction in the

BA/HS wage differential for standard demand and supply reasons. Thus, as is standard in examinations of technological change, the combination of movements in supply with movements in wages and wage ratios serve to identify the nature of movements in skill demand and, with it, the nature of technological change. In this case, what is needed for the period before 2010 is a type of technological change that would leave the skill wage ratio unchanged in the face of large increases in the supply of skill. This highlights the importance of clarity about the true pattern of movements in wages and wage ratios.

According to our model, the patterns in Figure 1 imply a period of technological progress induced by the education boom followed by a period that might fit with claims in other papers that technologically related productivity increases have failed to keep up with ongoing increases in education. But other depictions in which the BA/HS wage ratio declines over a longer time period would imply that the UK is experiencing a different technological/productivity path. We will examine wage and education group definitions, arguing that the pattern in Figure 1 is the right one for thinking about technological changes in the UK. And, perhaps more importantly, we will also argue that the declining ratio after 2010 arises from a source that suggests the decline does not reveal a decline in demand for skills relative to their supply.

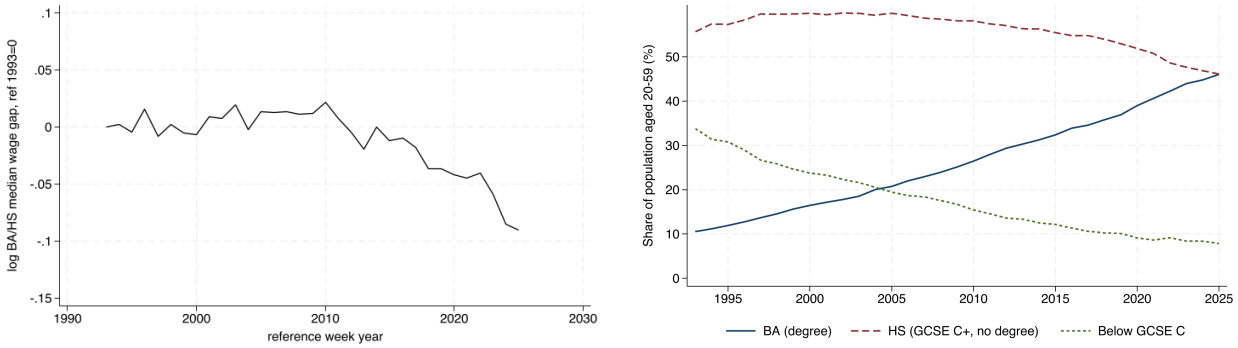


(a) Age-adjusted log BA/HS median wage ratio (b) Education-group population shares

Figure 1: The UK BA/HS wage ratio and the changing education composition, 1993–2025

*Notes:* LFS, persons aged 20–59, 1993–2025. Panel (a) is the age-adjusted log BA/HS median wage ratio (reference 1993 = 0), a replication and extension of Blundell et al. (2022) Figure 2; the hourly wage is last weekly pay divided by actual hours. Panel (b) is the share of the population aged 20–59 in each education group (BA, HS, and below GCSE C).

## 3.2 Regional Trends

The trend in figure 1 hides an important heterogeneity between London and the rest of the UK (Stansbury et al., 2023). In figure 2, we recreate our age adjusted BA-HS premia for

London and a combination of all the non-London UK regions. Here, we can see that the flat trend was common to London and the other regions up to 2010 but then they part ways. The trend remains flat for London but declines sharply for other regions. This converts the question of why the flat pre-2010 trend changed to a strongly declining trend after 2010 for the UK as a whole into one about regional differences in wage movements that coincide with concerns about the relative movements in productivity inside and outside London.

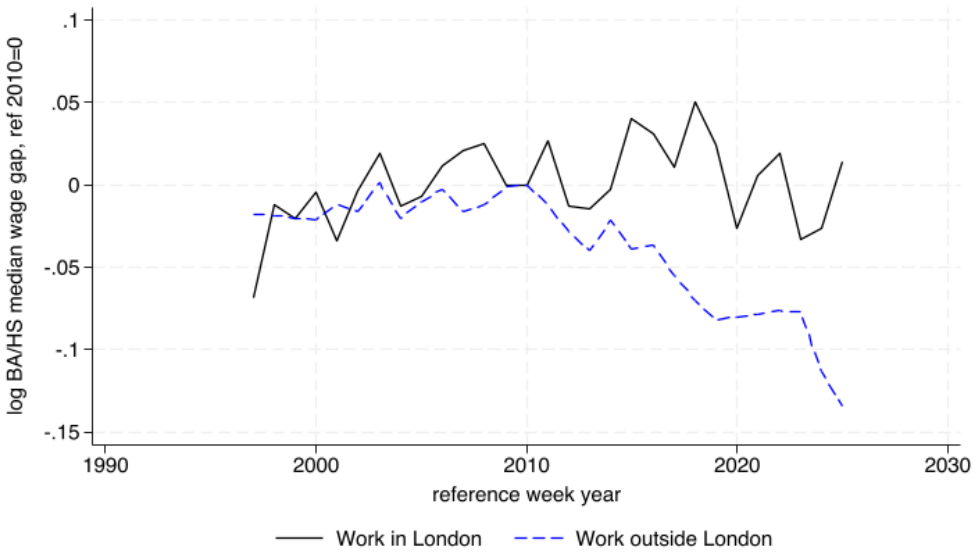


Figure 2: Age-adjusted log BA/HS median wage ratio, London versus rest of the UK, 1997–2025

*Notes:* LFS, persons aged 20–59, 1997–2025, by region of workplace; the hourly wage is last weekly pay divided by actual hours. The line is the age-adjusted log BA/HS median wage ratio, reference 2010 = 0. The solid line is London and the dashed line the pooled rest of the UK.

The figures so far are based on using all employees, but there are two reasons to focus our attention on private sector workers. First, pay in the public sector is set largely through national pay scales and pay review bodies, with limited regional differentiation beyond higher wages for London, and over much of the 2010s it was subject to a pay freeze followed by a 1% cap on annual awards. Public-sector relative wages therefore mainly reflect pay policy rather than local labour market conditions. Second, part of our interest is in the impact of technological change on production, which is more naturally studied in the private sector. In Figure 3, we recreate the regional educational wage differential series for private sector workers. The general pattern of flatness throughout the period for London workers, and before 2010 for non-London workers followed by a decline thereafter, continues to be present for private sector workers only. In levels, the age-adjusted private-sector premium outside London (the equal-weighted average across the eight age bands of the log BA/HS median ratio, whose changes are exactly the plotted series) was 0.439 log points in 2010, so that BA

median wages were 55% above HS median wages, and 0.285 (33%) in 2025. In London it was 0.383 (47%) in 2010 and 0.352 (42%) in 2025, but the London series is noisy from year to year and fluctuates around its 1997–2025 average of 45% with no clear trend. In the public sector (Figure 21 in Appendix A), the differential outside London also falls after 2010, but by much less: by about 2 log points by 2019 and 6.5 log points by 2025, against 10 and 15.5 log points in the private sector. The London public-sector series is too noisy to read a trend from, because the London public-sector HS cell is small. Nor is the private-sector pattern driven by changes in the gender mix within education groups: re-estimating the series on year × age band × sex cells with age-band × sex dummies leaves it unchanged, with a non-London decline over 2010–2025 of 13.5 rather than 15.5 log points and a flat London series (Figure 19 in Appendix A). The LFS results in the remainder of the paper are for private sector workers only; the ASHE data do not distinguish sectors.

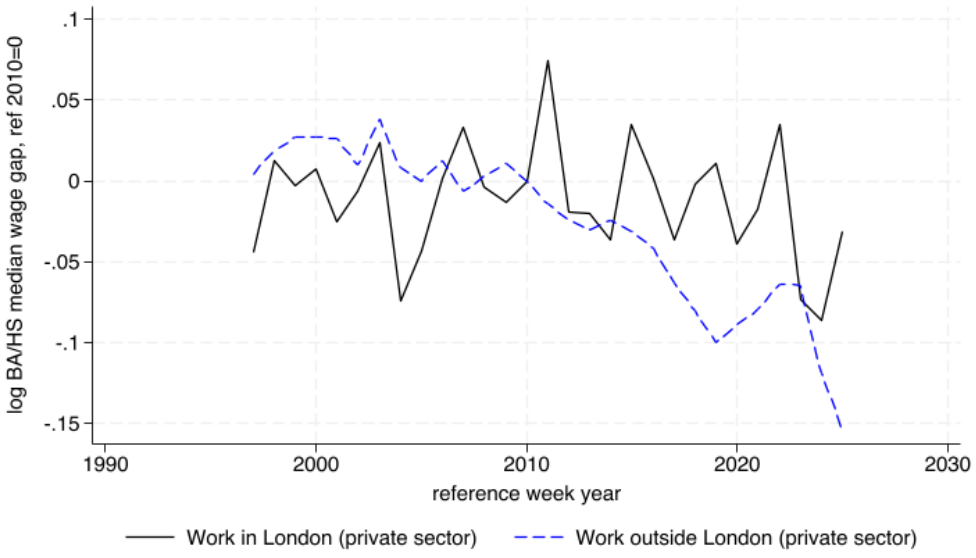


Figure 3: Private sector: age-adjusted log BA/HS median wage ratio, London versus rest of the UK, 1997–2025

*Notes:* LFS private-sector employees, persons aged 20–59, 1997–2025, by region of workplace; the hourly wage is last weekly pay divided by actual hours. The line is the age-adjusted log BA/HS median wage ratio, reference 2010 = 0. The solid line is London and the dashed line the pooled rest of the UK.

# 4 The Importance of Defining the High School Education Group

Before investigating potential explanations for the divergent trends in the BA-HS gap after 2010, it is important to reconcile the trend in figure 1 with other possible characterisations of

movements in the education wage differential for the UK. As described in section 2.2, there are multiple plausible definitions of lower education groups that could be compared to what we called BA workers. In figure 4, we plot BA-HS wage differentials over time for private-sector workers in all the non-London regions pooled, using three different HS definitions: 1) GCSE with a grade of at least a C but no university degree (including those who get A levels and/or tertiary education that does not result in a university degree) (from Blundell et al., 2022); 2) A levels only (i.e., not including people with GCSE exams of C or more who do not get A levels or people with tertiary education that does not result in a university degree) (from Stansbury et al., 2023);[7] and 3) any education below a BA (i.e., those without good GCSEs, with less than a C on the GCSE, along with those with a C or more on the GCSE's, A-levels and tertiary education that does not result in a university degree).

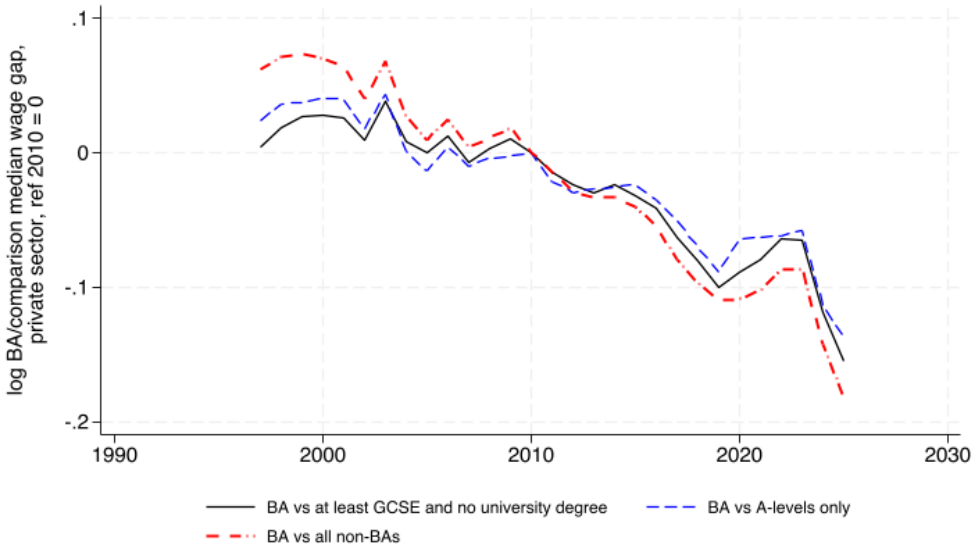


Figure 4: Age-adjusted BA–HS wage differential outside London for private-sector workers under alternative high-school definitions, 1997–2025

*Notes:* LFS private-sector employees, non-London regions (by region of workplace) pooled, persons aged 20–59, actual-hours hourly wage. Each line is the age-adjusted log ratio of the BA median wage to the comparison-group median wage, estimated as for Figures 2 and 3 (Section 2.2); reference 2010 = 0. Each line uses a different high-school definition: (1) GCSE grade C or above without a degree, including A-levels and sub-degree tertiary (Blundell et al. 2022); (2) A-levels only (Stansbury et al. 2023); and (3) any qualification below a degree, including those below GCSE grade C. The BA group is a university degree or higher throughout.

[7]Both BGJ and Stansbury et al. (2023) adjust for changes in age composition over time. Blundell et al. (2022) regress the log ratio of BA to HS median wages, computed within year by age-band cells, on year and age-band dummies, as described in section 2.2, while Stansbury et al. (2023) use a Mincerian log wage regression approach. In Appendix A, we show that age adjusting using the Mincerian regression results does not alter our main patterns. The appendix also contains the equivalent of Figure 4 for London, showing that there, too, the HS definitions other than the one including those without good GCSEs produce very similar BA premium trends.

The key takeaway from the figure is that the first two HS definitions generate the same BA-HS pattern over time, with a largely flat trend up to 2010 and a decline thereafter.[8] The definition including those without good GCSEs differs from the other two in that it declines before 2010 and declines faster than our main specification after 2010. These patterns are interesting for a focus on wage inequality, helping to answer the question of whether BA workers are pulling away from other workers in their wages. The answer outside of London over an extended period of time is that they are not. We argue later that the minimum wage plays a substantial role in these patterns and will have its greatest impact on the wages for a group that includes those without good GCSEs, potentially explaining why the differential relative to a median wage that includes that lowest group declines over a longer period. Given our desire to focus on the type of skill price differentials relevant for assessing models of technological change, wage premia that do not include wages driven directly by policy seem preferable. Based on that argument and the arguments in section 2.2 about other HS definitions, in the remaining sections of the paper, we work with the HS definition that uses GCSE with a grade of C or higher as the threshold for HS.[9]

# 5 Wage Patterns for Individual Education Levels

The time patterns of wages for the individual education groups that underlie the movements in the wage differential can provide extra information that is relevant for discerning among different theories of what is driving the differential. The standard supply-driven explanation for the decline in the differential after 2010 outside of London, for example, implies that we should observe that it is being driven by a decline in the BA wage and an increase in the HS wage.

In Figure 5, we present the wage series for BA and HS workers for London and the pooled non-London regions separately. We normalise all the series to 0 in 2010 to make it easier to see differences in time trends. Before 2010, the four series all move strongly together, fitting with the earlier observations of a lack of movement in the BA-HS differential inside or outside London. All the series show a strong (roughly 15%) increase from 1997 through to about 2002 then flatness until 2009. All the series also show a common, substantial drop in the aftermath of the 2008 recession. The London HS and BA series recover very little in

[8] Stansbury et al. (2023) present education wage premium graphs by region for the UK from 1997 through 2019. Our estimates agree with theirs in the finding that the BA-HS gap is flat in London throughout, and outside London it falls substantially over 1997–2019 taken as a whole. They differ on when within that window the fall occurs; Table 9 in Appendix A.4 decomposes the difference.

[9] Appendix A contains the same plot as in Figure 4 for London, showing that the BA differential is flat for London across the whole time period regardless of the HS definition, the one qualification being that the A-levels-only line runs below the others through much of the period and ends furthest below.

the ensuing decade and a half. Interestingly, BA wages outside London follow very much the same pattern, moving quite closely with the BA wages inside London. But what stands out most is the nearly 12% rise in the HS wage for the non-London regions between 2014 and 2025, on a nearly continuous upward trend. Thus, the decline in the BA-HS premium outside London is driven primarily by an increase in HS wages rather than by declines in the BA wage outside London relative to in London.[10]

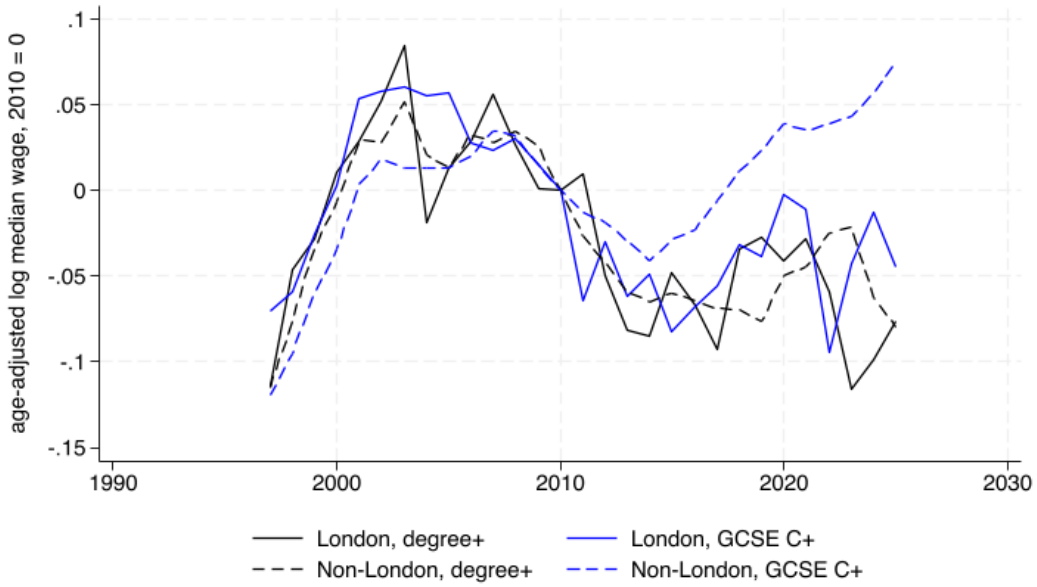


Figure 5: Age-adjusted log median real hourly wage by education group, London versus non-London, private sector, 1997–2025

*Notes:* LFS private-sector employees, aged 20–59, 1997–2025, by region of workplace, age adjusted. The hourly wage is measured using actual hours.

At this point, we see education levels and wages in the UK over the last three decades as characterised by a set of four patterns. Prior to 2010, the education level of the workforce grew strongly but the BA-HS wage differential was unchanged both in London and in regions outside London as the wage levels moved strongly together in all education x region groups. In Blundell et al. (2022), we argue that this pattern can be explained by an endogenous technological choice model in which firms respond to the educational increases by shifting to a technological frontier defined in the US, and we provide evidence in support of the model using data from various sources. The second pattern is the coincident movement of HS and BA wages in London after 2010, implying a lack of change in the BA-HS differential and continuing the pre-2010 path. The third pattern is the coincident movement of BA wages in

[10]The same exercise run on public-sector employees shows the four series continuing to move together after 2010, with no comparable break-out of the non-London HS wage (Figure 22 in Appendix A).

London and in the non-London regions, also continuing the pre-2010 path. The fourth is the rise in real wages for HS workers outside of London after 2010 relative to both HS wages in London and BA wages both inside and outside London. We are interested in understanding the source of the post-2010 wage patterns, noting that they have arisen in the presence of ongoing increases in education levels across the various UK regions. In the remaining sections of the paper, we consider candidate explanations for the post-2010 patterns with the goal of narrowing the set of candidate models that could explain this complicated array of wage and employment movements.

# 6 Relative Supply Effects and the post-2010 Patterns

The most straightforward explanation for the post-2010 decline in the university wage premium is that ongoing increases in the supply of workers with a BA have outstripped demand. The regional differences highlighted in Stansbury et al. (2023) and Xu (2023) add one qualifier to this hypothesis: BA supply would have to have outstripped demand outside of London but been exactly met with a relative demand increase in London. In Figure 6a, we plot three year averages of the log of the BA/HS wage premium for each of the 12 regions against the average share of BA's among private-sector wage earners in the same time x region cell. We do this only for 30 to 39 year olds in order to hold age constant - something we argued earlier is important for obtaining the true education premia. The hollow dots correspond to London. Because the graduate share increases with time, moving rightward along the horizontal axis also corresponds roughly to time. As we have seen before, with time, in London, the BA share rises but the education wage premium is essentially unchanged. In contrast, the solid dots, corresponding to other regions, show a clear negative slope (though, note that for the other regions, time and BA share do not match as closely because the high versus low share dots capture both persistent differences across regions and differences over time within regions).

Figure 6a is an updating in somewhat different form of Figure 7 in Stansbury et al. (2023) and serves to make the same point: for the regions outside of London, higher BA shares are associated with lower BA/HS wage premia (the simple correlation between the regional shares and premia is $-0.71$). While this simple correlation is powerful, we can investigate the hypothesis that excess supply is driving down the premium further by working with a standard panel model in which we control for region and common time effects. This serves to answer the question: are increases in the BA share in a region relative to UK-wide trends associated with declines in the wage premium, also measured within the region and relative to common trends? It is important to move to this specification because both the BA share

and the BA/HS premium outside London are trending over time so any association between them could reflect anything driving those trends - not just a demand/supply driven link. We implement the panel estimator by separately regressing the three year averages of the BA/HS wage premium and the BA share for each region on a complete set of region and time fixed effects. In Figure 6b, we plot the residuals for each region x time cell from the BA/HS premium regression against the matching residuals from the BA share regression. We again restrict our sample to 30 to 39 year olds. Here, the picture is very different: there is no association between the wage premium and the BA share residuals (the correlation between the share and the premia is now +0.17).[11] In other words, it is not the case that the regions in which the BA share grew the most are the ones where the BA/HS premium fell the most. We view this as strong evidence against relative supply as a key driving force for the movements in the BA/HS wage premium.

These patterns are correlations, and the regional graduate share should not be treated as exogenous. Graduates are geographically mobile and tend to move towards regions offering better economic opportunities (Xu, 2025). Region fixed effects remove persistent differences in opportunities, but not region-specific changes over time: a region experiencing a rise in demand for graduates would tend to attract graduates at the same time as its premium held up, and this would bias the residualised correlation upwards and away from the negative relationship predicted by the supply hypothesis. Causal evidence therefore needs a source of variation in graduate supply that is not driven by local demand. Blundell et al. (2022) provide such evidence for 1993–2016. They use shift-share instruments that exploit regions' initial differences in cohort composition and in the education of the parents' generation interacted with national growth in graduate share. They find that increases in the relative supply of graduates had close to a zero effect on the graduate wage premium and on the level of the graduate wage. Together with the patterns in Figure 6b, this makes it difficult to attribute the post-2010 decline in the non-London premium to graduate supply outstripping demand.

An associated argument relating to the excess supply hypothesis is that the BA/HS premium is lower outside of London because BA educated workers outside of London are more likely to have to take jobs that don't actually require a BA, i.e., they are being paid the HS wage. Being pushed into non-BA jobs is a direct reflection of there being insufficient demand for BA workers. Xu (2023) shows that it is indeed the case that BA educated workers working outside London are less likely to work in jobs that require a university education than those in London. Working with the LFS, Xu (2023) finds that 65% of BA workers in

[11] The correlation between the share and the premium in London is positive and sizeable in both specifications. To match the frequency being used for the other regions, we can work in first differences. In that case, the correlation for London is −0.39.

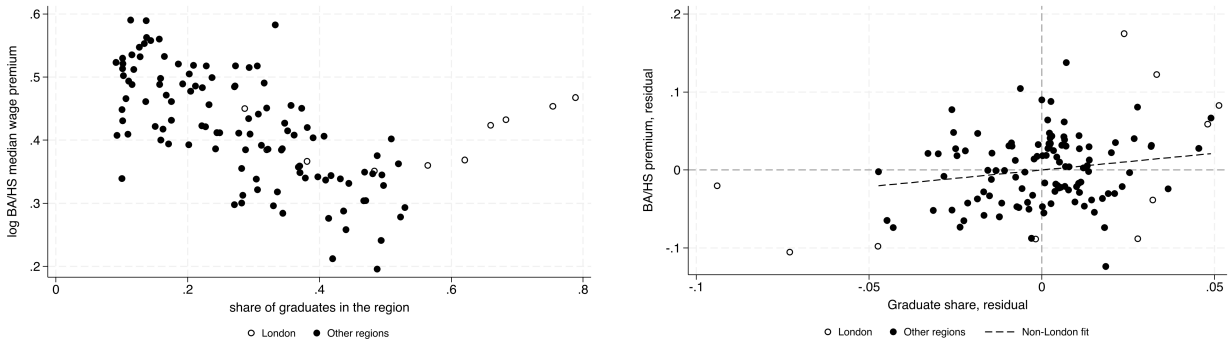


(a) Unconditional (b) Residualised Relative to Region and Year FE

Figure 6: BA share and the BA/HS median wage premium, private sector, 30–39 year olds only, by region and period, 1997–2025

*Notes:* LFS private-sector wage-earners aged 30–39, by region of workplace, 1997–2025. Each point is one region × 3-year-period cell; hollow markers are London, solid markers the other eleven regions. The vertical axis is the cell's log BA/HS median wage premium (actual-hours wage); the horizontal axis is the cell's graduate share. The cross-cell correlation between the premium and the graduate share is −0.71 outside London (110 cells) and +0.42 within London (10 cells) in the unconditional graph. In panel (b) each of the premium and the share is regressed on a complete set of region and 3-year-period fixed effects and the residuals are plotted against one another; there the correlations are +0.17 outside London and +0.54 within London.

London in 2022 were in occupations that required a BA while in Lincolnshire that number was only 42%. Moreover, the percentage declined in some regions outside London but edged up in London. As Xu points out, these differing trends have the potential to explain the declining BA-HS differential outside versus in London.

In order to assess the importance of the under-placement mechanism identified by Xu, we implement a panel data specification in which we ask whether differential movements in wages outside versus in London after 2010 can be explained by differential movements in the share of BA workers who are working at lower qualification jobs. As a first step, we need a measure of whether a given occupation requires a university level education. Following Xu (2023), we use the mapping in Home Office and Immigration (2014) to map the SOC occupation categories in the LFS into National Qualification Framework (NQF) levels that are used by the Home Office as part of the immigration system.[12] There are four potential NQF groupings that are relevant for our exercise: Levels 6-8 (occupations requiring a BA or higher education); Levels 3-5 (occupations associated with A Levels and/or some post-secondary education less than a BA); Level 2 (occupations requiring a GCSE grade of C or greater); and Level 1 (corresponding to the skills of those without good GCSEs).

[12]NQF was replaced by the Regulated Qualifications Framework in 2015. They share the same levels, from entry level up to Level 8, so the shift is not consequential for our exercise. We maintain the NQF acronym throughout.

Unfortunately, the Home Office mapping between occupations and NQF levels groups levels 1 and 2 together even though we would prefer to group level 2 with level 3, and, so, we end up with three skill groupings. We will call Level 6+ "high skilled" jobs, Levels 3-5 "middle skilled", and Levels 1-2, "low skilled".

In Figure 7, we plot the share of private-sector BA educated workers in the LFS who are in high skilled and middle skilled jobs in and outside London. The figure replicates Xu (2023)'s result in our data. In particular, the shares are substantially lower outside of London and show some trend down over time outside London. Importantly, though, there is little downward trend outside London after 2010.[13] Table 1 reports how each education group in the private sector is distributed across the three NQF occupation bands over the years in which the SOC2010 occupational system was in use in the LFS (2011–2020). Consistent with Figure 7, 53.4% of BA workers are employed in NQF Level 6+ occupations, against 17.7% of HS workers and 5.0% of those below GCSE grade C. Among HS workers, 35.0% are in the middle-skilled Level 3–5 band and 47.3% in low-skilled occupations below Level 3, so nearly half of HS workers are in occupations below the middle-skilled threshold. Note that even among BA workers, 19.5% are in below-Level-3 occupations.

Table 1: Share of each education group working in each NQF occupation band, private-sector employees, LFS 2011–2020

| Education group | High (NQF 6+) | Mid (NQF 3–5) | Low (<NQF 3) |
|---|---|---|---|
| BA | 53.4 | 27.1 | 19.5 |
| HS | 17.7 | 35.0 | 47.3 |
| Without good GCSEs | 5.0 | 22.9 | 72.1 |

*Notes:* Each cell is the weighted within-education share of employment in the NQF occupation band, so rows sum to 100. LFS 2011–2020 (the years coded to SOC2010); private-sector employees aged 20–59 reporting a main-job occupation; person weights. NQF bands map 4-digit SOC2010 unit groups to skill levels via the Home Office 2014 Codes of Practice (Home Office and Immigration, 2014). "BA" = degree or higher; "HS" = GCSE grade C or above without a degree; "Without good GCSEs" = below GCSE grade C.

Using these classifications, we run regressions with the dependent variable being the age-adjusted log median wage of private-sector employees at the level of region X education group X year, using LFS data from 2001 to 2020, where the regions are London plus the 11 other official first-level statistical regions of the UK, defined in the Data section. As covariates, we include region and year fixed effects as well as an interaction of the London dummy variable with a post-2010 dummy variable to capture the changes in the relationship

[13]Because of changes in occupation classification every decade or so, each of the different segments of the plotted lines should be looked at independently. Those show declines in the share in NQF 6+ jobs for BA's outside London in both of the first two segments.

Figure 7: Graduates' share working in NQF Level 6+ and Level 3–5 occupations, London versus non-London, private sector

*Notes:* LFS, 2001–2025; private-sector employees reporting a main-job occupation, aged 20–59, by region of workplace; person weights. Lines show the share of graduates (BA) working in NQF Level 6+ and in NQF Level 3–5 occupations, where NQF bands map 4-digit SOC unit groups to skill levels via the Home Office 2014 Codes of Practice. The line segments break at the SOC reclassifications (2010/11 and 2020/21) and should be read separately.

between London and non-London wages after 2010 for the various education groups. Our interest is in that interaction, which represents the difference between the median wage in London and the average of those in all other regions combined after versus before 2010. We then add in variables corresponding to the proportion of workers in the given education x region x year cell in NQF level 6+ occupations and the proportion in level 3 to 5 occupations (leaving the proportion in level 1 and 2 occupations as the base). If differential access to higher-skilled occupations is determining the change in the London-nonLondon differentials before versus after 2010 then the estimated coefficient on the London-Post2010 interaction should move to zero.

Table 2: Age-adjusted log median wage: London post-2010 differential with NQF occupation controls, private-sector employees

| | Graduates | | High school | |
|---|---|---|---|---|
| | Base | +NQF | Base | +NQF |
| London × post-2010 | 0.005 | -0.017* | -0.071*** | -0.068*** |
| | (0.007) | (0.009) | (0.005) | (0.006) |
| NQF high-skill share | | 0.712*** | | 0.328 |
| | | (0.094) | | (0.375) |
| NQF mid-skill share | | 0.513** | | 0.060 |
| | | (0.233) | | (0.183) |

Standard errors in parentheses

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

*Notes:* Private-sector employees aged 20–59, 2001–2020. Each column is a separate regression on the 240 region × year cells of one education group (graduates or HS); the hourly wage uses actual hours. All regressions include region and year fixed effects, are unweighted, and cluster standard errors by region (12 clusters); $N = 240$ in every column. Column "Base" includes *London × post-2010* only (London and post-2010 main effects absorbed by fixed effects). Column "+NQF" additionally controls for the share of the education–region–year cell's workers in NQF Level 6+ occupations and the share in NQF Level 3–5 occupations (Level 1–2 omitted), measured on the same private-sector 20–59 sample as the dependent variable. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

We report the results from this exercise in table 2. The first two columns correspond to the BA education workers. The London × post-2010 coefficient is 0.005 (standard error 0.007): once wages are age-adjusted, BA median wages in London show no significant change relative to the average across non-London regions after versus before 2010, in line with the close co-movement of the two BA series in Figure 5. In the second column, we add the NQF proportions. They are strongly associated with graduate wages: relative to the share in Level 1–2 occupations, the coefficient on the share in NQF Level 6+ occupations is 0.712 (0.094), so a ten-percentage-point higher share goes with a roughly 7% higher graduate median wage, and the coefficient on the Level 3–5 share is 0.513 (0.233). The occupational placement emphasised by Xu (2023) is therefore related to graduate wage levels across regions

and years. But there is no relative post-2010 London graduate wage gain for it to explain. Adding the NQF shares moves the London × post-2010 coefficient to −0.017 (0.009). This means that conditional on occupation, the non-London graduate wage increased by about 2 percent in relative terms after versus before 2010. This is offset by the relative deterioration in non-London graduates' occupational placement seen in Figure 7, so that overall there was not much difference in the graduate wage trend in versus outside London.

In the second set of columns of table 2, we see that the relative median wage for HS workers in versus outside London fell by 7.1% after 2010 (−0.071, standard error 0.005), which fits with our observation that it is increases in HS wages outside London that are driving the relative decline in the BA-HS differential outside London.[14] Of the 0.076 log-point relative movement in the BA-HS differential implied by the two base columns, the BA column accounts for only 0.005, less than a tenth, and that part is not statistically distinguishable from zero.[15] Thus, the explanation for the movements in the BA-HS differential lies almost entirely with the rise in HS wages outside of London. This fact serves to under-cut the relative supply explanation more generally since it is unlikely that an over-supply of BAs outside London would show up only through an increase in HS wages, with little or no effect on the BA wage.

# 7 The Post-2010 Rise in the Non-London High School Wage

We turn now to investigating what we view as the main puzzle arising out of our data so far: why, in an era commonly viewed as being characterised by skill-biased technological change, did the real wage for HS workers outside of London rise, doing so while the HS wage in London was declining. In this section, we investigate a set of compositional and other direct factors that might explain this rise.

## 7.1 Measurement Issues

The most obvious explanation for the striking departure of HS wages from expectations is that it is not actually true. The UK LFS, in recent years, has faced issues with response

[14]For workers below the HS level, the London × post-2010 estimate is more negative than for the HS group. These results are available from the authors.

[15]It is interesting to note that adding in the NQF shares has almost no impact on the HS relative wage differential (the London × post-2010 coefficient moves from −0.071 to −0.068), indicating that it is not altered access to particular skill level jobs that is driving the relative increase in HS wages outside London. We return to this point in section 8, when we examine compositional explanations for HS wage patterns.

rates and some irregularities. According to the ONS Labour Force Survey quality monitoring report (Office for National Statistics, 2026c), response rates fell from just over 50% for new survey entrants in 2019 to rates around 30% by 2021 and remain in the latter range till 2025Q4. In Appendix B, we compare elements of the wage distribution as depicted in the LFS with those from ASHE. Since ASHE is based on administrative data, it does not face the response rate issues in the LFS. Examining comparable wage measures for workers aged 25 and over, we find that the median wages found in LFS and ASHE data are very similar in 2015, 2019 and 2025. However, the two datasets show noticeable discrepancies in the lower tail in more recent years. In 2025, the 10th percentile of the wage distribution for elementary occupations was £11.5 in ASHE but only £9.1 in the LFS. Higher recorded wages for workers paid a fixed hourly wage in the LFS suggest that part of the problem may be related to hours reporting since the majority of even low wage workers have wages derived using earnings and hours measures.

So, could the HS wage patterns in the LFS be driven by measurement problems? If hours are overstated uniformly across groups, and increasingly so over time, then all four lines in Figure 5 drift down together, and the finding of the non-London HS wage growing much faster than the other three is unaffected. If instead hours are increasingly overstated among lower-paid workers with variable hours, who are more likely to be HS workers outside London, then non-London HS wage growth is understated in Figure 5. Thus, our reading of Figure 5 is robust to this form of bias; if anything, it is conservative. There is not enough evidence to indicate how non-response to the LFS should bias the four wage trends in Figure 5. A wage trend is biased only if, within an education group and region, the response rate of higher- relative to lower-paid workers has shifted over time; and only if that shift differs between London and the rest of the UK would it affect the finding that the non-London HS wage grew the fastest. The ONS linkage analysis reports the LFS shortfall by age and sex, but it does not report it by pay level (Office for National Statistics, 2026b).

We cannot compare the LFS wages directly with ASHE figures because ASHE does not include an education variable. But we can draw some inferences that suggest the LFS pattern is genuine. First, ASHE and LFS data have similar wage values at the 25th percentile and above - a range that would include the median HS wage that we are using for our LFS figures. Second, we can use occupation groupings according to the NQF levels described in the previous section to examine wage movements in high and middle skilled occupations in both the LFS and ASHE to see if they are comparable. Recall that high skilled occupations are ones that require a BA or more and, so, could be expected to have wage movements that are similar to those for BA workers. The middle skilled group corresponds to A-levels plus other education greater than high school but less than a university degree. In that sense, it is similar in nature to one definition of HS wages. In contrast to our earlier plots, which

work with median wages, the nature of the ASHE data implies we need to work with mean wages for this exercise.[16] Patterns in the LFS are similar between mean and median wages.

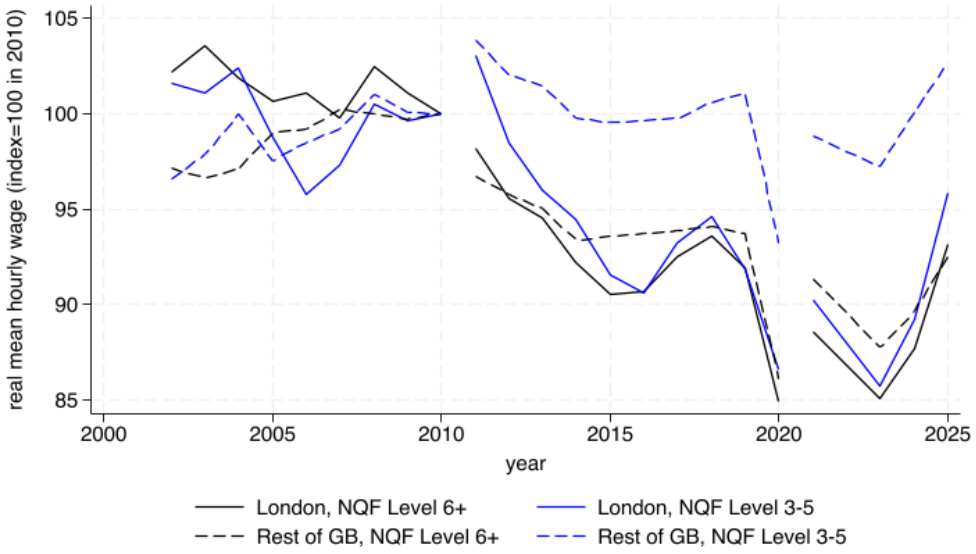


Figure 8: Real hourly wage trend by NQF occupation band, London versus rest of GB (ASHE, 2002–2025)

*Notes:* ASHE Table 15.5a (4-digit occupation × region, region of workplace; Great Britain, so "rest of GB" excludes Northern Ireland). Within each NQF band the series is the ASHE employment-weighted mean of the published occupation×region mean hourly wages, by London and non-London, deflated to 2025 prices and indexed to 100 in 2010. Each line is drawn as three separate segments, breaking at the SOC reclassifications (2010/11 and 2020/21), and should be read segment by segment; the 2010 base is common to all three, as in the LFS counterpart of Figure 23. Cells with suppressed mean wage or employment are excluded; the employment share of suppressed cells in London ranged from 14.8% to 19.1% in 2020–2022, against 2.0–6.6% over the 2010s. The 2022 points are dropped because of that suppression and the implausible recorded values in London discussed in the text.

Figure 8 plots the employment-weighted mean hourly wages for each skill group, aggregated from published ASHE tables by occupation×region. The averages are computed separately for London and the rest of the country, deflated to 2025 prices and indexed to 100 in 2010.[17] The Figure shows a pattern that is somewhat similar to what we observed

[16]In particular, we work with published ASHE statistics at the 4 digit occupation and region level (Table 15.5a), which means we have mean and median wages by occupation and cannot construct the median wage for our NQF groups. We omit 2022 from the figure due to data quality concerns. The ASHE tables suppress small cells. The employment share of suppressed cells in London ran at 14.8–19.1% over 2020–2022, compared to 2.0–6.6% over the 2010s. Among the cells we do observe, occupations accounting for about 53% of London employment recorded a 2022 mean below 90% of *both* their 2021 and 2023 values—a dip-and-rebound that appears implausible (for example, financial managers: mean pay £48 → £33 → £56 and recorded employment 134,000 → 49,000 → 147,000 over 2021–23). We therefore drop the 2022 points rather than let this artefact distort the series.

[17]We present the corresponding figure using NQF groupings in LFS data in Appendix A (Figure 23). There, too, the middle skilled wage outside of London is an outlier, lying well above the middle skilled wage in London.

for education group wages in the LFS: since 2010, the "London, NQF Level 6+", "Rest of GB, NQF Level 6+" and "London, NQF Level 3–5" series have all declined substantially in real terms and closely together, while the "Rest of GB, NQF Level 3–5" series stands out as having flatlined. Thus, as in the LFS education data, the high-medium skilled wage differential is unchanged over time in London but declines after 2010 outside London, and the latter decline is driven by better performance in middle skilled wages outside versus inside London.

Finally, it is worth mentioning that a third source of data – personal tax and education records – also yields patterns consistent with what we have seen so far. Using linked tax and education records (LEO) for successive GCSE cohorts in England, Figure L.1 in Britton et al. (2026) shows that median real earnings of non-graduates below age 30 grew substantially faster across recent cohorts than those of young graduates, for both men and women. Again, this indicates a decline in the graduate wage premium driven by a strong rise in HS type wages – in this case from a source free of survey non-response.[18]

Taken together, these comparisons imply that the patterns observed in the LFS are genuine. In the analysis that follows, we draw on both the LFS and ASHE.

## 7.2 The Minimum Wage

In real terms, the UK minimum wage for adults was broadly flat between 2007 and 2015, then it increased by 6% between 2015 and 2016, and then by a further 24% between 2016 and 2025 (Figure 9). Recall that the rise in the non-London HS wage also occurred around the mid-2010s. Thus, a second possible explanation for this rise is that it was driven by spillovers from the minimum wage as employers sought to maintain wage differentials, possibly for fairness related efficiency wage reasons. A set of recent papers examine wage spillover effects from the minimum wage in different economies: Cengiz et al. (2019) and Gopalan et al. (2021) for the US, Brochu et al. (2025) for Canada, and Giupponi et al. (2024) for the UK. All find evidence of spillovers, with the spillovers occurring up to about $2 to $2.50 above the minimum wage in Canada and the US, and about £2 above the minimum wage in the UK. While these are all in different currencies, they all amount to spillovers reaching about 20%

[18]Britton et al. (2026) focus part of their discussion on a different secondary education group, one composed of individuals who achieved at least the points equivalent of five grade C GCSEs in their best eight GCSE subjects and who have a Key Stage 5 record (that is, entered A-level or equivalent Level 3 study) within three years of sitting their GCSEs, limiting the analysis to people "who could have chosen to enter university at, or soon after, age 18". This group is more positively selected than our HS group, which includes everyone with at least a GCSE grade C (or equivalent) but no degree, whether or not they continued to Level 3 study. Working with that definition, they find much more similar earnings trajectories across cohorts for university graduates and non-graduates (their Figure L.3).

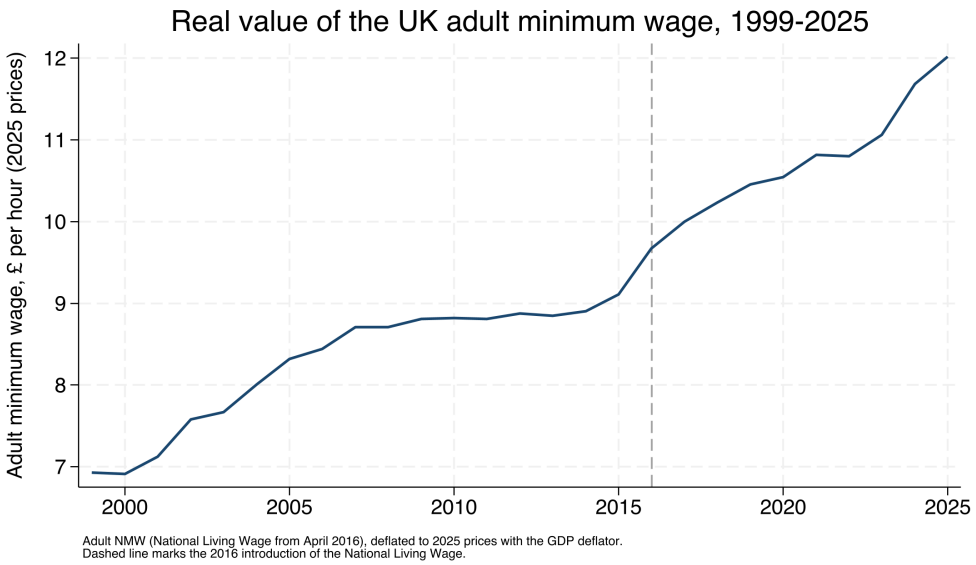


Figure 9: Real value of the UK adult minimum wage, 1999–2025

*Notes:* The UK adult minimum wage—the National Minimum Wage, and the National Living Wage from April 2016—for 1999–2025, in 2025 prices (deflated with the GDP deflator). The dashed vertical line marks the April 2016 introduction of the National Living Wage.

above the minimum wage, and Brochu et al. (2025) show that the spillover effects are about relative comparisons, with the spillover patterns being invariant to the level of the minimum wage so there is reason to think that these papers are delivering a common lesson in terms of the extent of spillovers from the minimum wage. We should be clear at the outset that the reach we find below is longer than this. The median HS wage outside London still sits some £3 above the NMW in 2025, and the occupation-level evidence in section 7.2.2 shows wage growth extending further still. Our results are therefore not simply a confirmation of the existing estimates. They suggest that when the floor rises as far and as fast as it did in the UK after 2015, the spillover reaches higher into the distribution than the earlier literature, estimated on smaller increases, had occasion to detect.

### 7.2.1 Regional Variation

Given these spillover patterns, is it plausible that the minimum wage could have affected median HS wages in the UK? In Figure 10, we plot the number of regions that lie in four gap bins defined by the value of the median HS wage in the region minus the national minimum wage (NMW): a difference of under £1; a £1 to £2 difference; a £2 to £3 difference; and a difference of over £3. The differences are calculated in nominal (contemporaneous) terms. Before the 2016 increase the NMW sat well below every regional median: nine of the twelve regions were more than £3 above it in 2010 and seven in 2015, and none was within £2. The

2016 rise meant only 3 regions remained in the over £3 above band, and two entered the within £2 band. The counts then partly reverse, but this reflects inflation rather than the floor receding, since the bands are nominal and a £3 gap in 2025 is worth far less than the same gap in 2016.[19] We view the nominal wage patterns as relevant because any potential spillover effects could reflect fairness based concerns and those would be based on nominal wages. We present a counterpart to figure 10 using public sector wages in Appendix A. The public-sector version provides a useful contrast, since our mechanism implies that the floor should bite in the private sector and not in the public one: all twelve regional public-sector HS medians are more than £3 above the NMW in every year to 2016; at the most compressed point, 2023, nine regions are still in the £3+ band; and no region's median is ever within £1 of the floor.

Given the results in the literature on minimum wage spillovers, the NMW might have started to affect median HS wages in multiple regions after 2016. Moreover, the minimum wage movements have the potential to explain the differences in movements in the HS wage inside versus outside of London. In Figure 11, we plot the NMW, the mean of the median HS wages across the 11 non-London regions, and the median HS wage in London. We can see the NMW getting close to the median HS wages outside London with time after 2016. But while it moves toward the London median HS wage, even by 2025, the NMW is over £5 below the London wage. Thus, if there are spillover effects, we would expect that they could drive the non-London HS wage up without affecting the London HS wage.

We investigate the potential effect of minimum wages by creating a minimum wage bite variable defined as:

$$b^e_{rt} \;=\; \frac{\text{NMW}_t}{\text{Median}^e_{r,2010}}. \tag{1}$$

where $\text{NMW}_t$ is the NMW in year t, and $\text{Median}^e_{r,2010}$ is the median wage in education group e and region r in the base year, 2010. This variable will be larger when the bite of the minimum wage is greater, i.e. when the NMW gets closer to the median wage in a region. Our approach is to estimate the effect of the bite variable on the median wage for each education group. Given that, we work with the median wage in a base year in order to reduce the endogeneity that would arise by definition if we used the current year's median wage in the denominator.

We implement our approach by estimating the following panel regression for each education group:

$$y^e_{rt} \;=\; \alpha_r + \delta_t + \beta\, b^e_{rt} + \varepsilon_{rt}, \tag{2}$$

[19] We report that real-terms version of this figure as Figure 25 in Appendix A: in 2025 prices, the number of regions more than £3 above the NMW falls from twelve in 2010 to nine in 2017 and five in 2025.

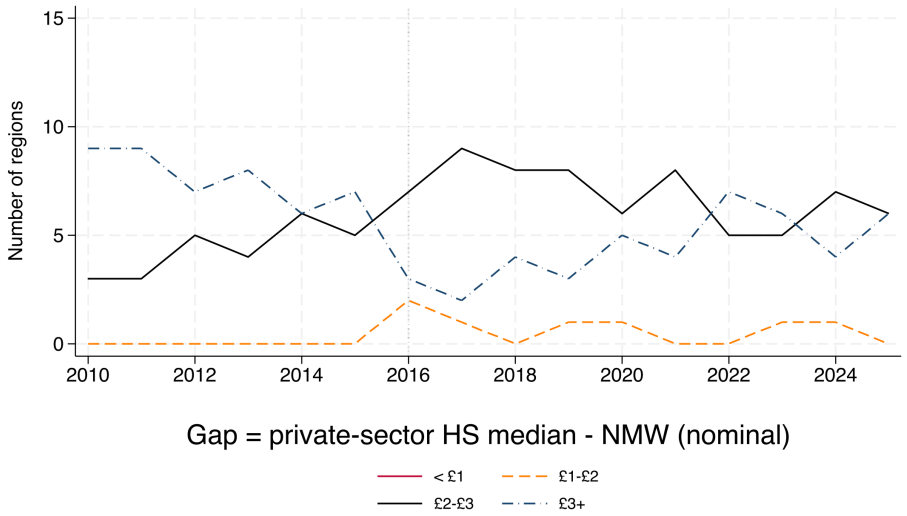


Figure 10: Number of regions by the nominal gap between the private-sector HS median wage and the minimum wage, by year

*Notes:* LFS HS wage medians by region for 20–59 year olds in the private sector, and the statutory NMW. For each year the twelve UK regions (including London) are grouped by the gap between the region's nominal HS median wage and the nominal NMW into four bands (<£1, £1–2, £2–3 and £3+); each line gives the number of regions in each band. HS = GCSE grade C or above without a degree.

where $y^e_{rt}$ is the log of the age-adjusted median wage for education group e in region r in year t. The observations used in the regression are region×year cells over 2010–2025, using all the years for which the necessary LFS variables are consistently coded. The regression includes region fixed effects, $\alpha_r$, and year fixed effects, $\delta_t$. The standard errors are clustered by region. All region-level regressions are weighted by the region×education cell's employment. The bite variable is arguably exogenous since its variation stems from national level minimum wage changes whose regional impacts are scaled by predetermined regional median values. The identifying variation for $\beta$ that this specification employs comes from comparing changes in regional median wages (relative to the overall trend) with movements in the NMW scaled by the regional base period median wage. A positive value for $\beta$, potentially indicating a spillover effect of the minimum wage onto the median wage, would arise if regional median wages moved with the NMW most in regions with lower values of the base median wage. This means that the identifying variation comes, ultimately, from differences across regions in the base median wage. Identification then relies on an assumption that changes in $\varepsilon_{rt}$ are independent of the base level of the median wage in the region. If they were not then $\beta$ would partially reflect whatever forces drive a relationship between the base median and median wage growth in a region. We provide an investigation of this assumption below.

Table 3 reports the results for our specification using HS wages. We present results for all workers and for the private sector only. Since London is something of an outlier in terms

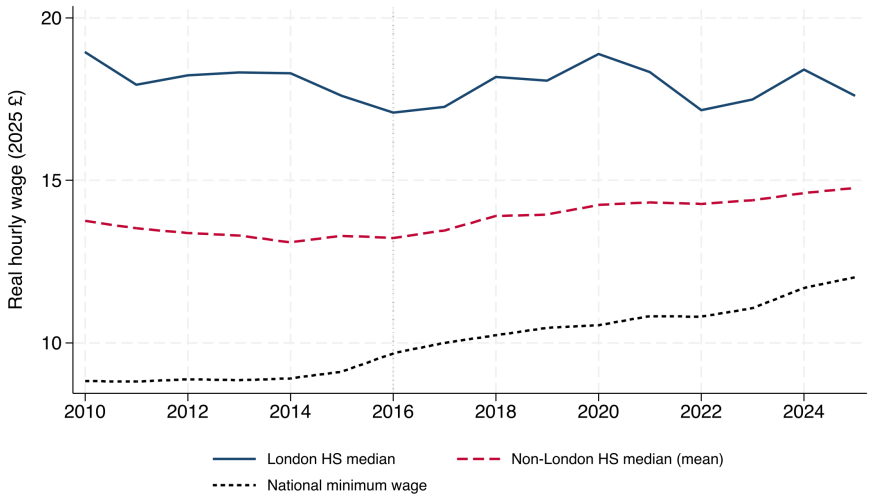


Figure 11: Real minimum wage, London HS median wage, and the mean non-London HS median wage, private sector, 2010–2025

*Notes:* LFS HS median wages for 20–59 year olds in the private sector, and the statutory NMW, 2010–2025, in 2025 prices. HS = GCSE C+ (EDU 2), by region of workplace. The non-London series is the unweighted mean of the eleven regional HS medians. The dotted vertical line marks the 2016 National Living Wage.

of its wage and bite variable values, it is important to check whether it is driving results. For that reason, we present results both with and without London. The results are striking. The $\beta$ coefficient is positive in every specification, but it is larger and more precisely estimated for the private sector. There it takes a value of 1.513 across all regions and 1.600 excluding London, in both cases statistically significant at the 1% level, so dropping London leaves the estimate essentially unchanged. For all workers, the coefficient is 1.202 across all regions, also significant at the 1% level, but it falls to 0.620 and is no longer statistically significant once London is excluded. The private-sector estimates are therefore both larger and more robust to the exclusion of London, and it is the private-sector coefficient that we carry into the counterfactual below.

To interpret the magnitude of the $\beta$ estimates note that it is a slope on the ratio NMW/Median, so dividing it by the median converts it to the effect of the minimum wage on the wage: $\partial y/\partial \text{NMW} = \beta/\text{Median}$. The real NMW rose from £8.82 to £12.02 over 2010–2025, an increase of £3.20. In the private sector—the sample plotted in Figure 11 and the one we carry into the counterfactual below—the 2010 HS median (in 2025 prices) was about £18.95 in London and £14.00 in the rest of the UK (the latter aggregated as the reciprocal of the employment-weighted mean of 1/Median). Applying the private-sector estimate $\beta = 1.513$, the implied effect $\beta \cdot (3.20)/\text{Median}$ is 0.26 log points at the London median and 0.35 at the non-London median (recall that the dependent variable, $y$, is actually the log of the median). These levels are not separately identified—the year effects absorb whatever

Table 3: HS wage on the minimum-to-median bite

| | All workers | | Private sector | |
|---|---|---|---|---|
| | All reg. | Excl. Ldn | All reg. | Excl. Ldn |
| $\text{NMW}_t/\text{Median}_{r,2010}$ | 1.202*** | 0.620 | 1.513*** | 1.600*** |
| | (0.221) | (0.376) | (0.144) | (0.227) |
| Observations | 192 | 176 | 192 | 176 |

Standard errors in parentheses

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

*Notes:* Outcome is the age-adjusted median HS log real wage year effect, region×year, 2010–2025. Region and year fixed effects; weighted by cell employment; standard errors clustered by region (12 clusters in the all-region columns, 11 in the columns excluding London).

the minimum wage did on average—but their difference is: the minimum wage raised HS wages outside London by about $\beta \cdot 3.20 \cdot (1/14.00 - 1/18.95) \approx 0.090$ log points *more* than in London over 2010–2025. In words, the rising floor lifted non-London HS wages about nine percent relative to London's, simply because the same national increase was a larger share of the lower non-London median. (The all-worker estimate gives a smaller differential: with $\beta = 1.202$ and 2010 medians of £19.43 and £14.60, the implied effects are 0.20 and 0.26 log points and the differential 0.065.) If the identifying assumptions are true, differential spillover effects from the rising NMW had a substantial effect on the difference in HS wage movements inside versus outside London after 2015.

In Table 4, we present the results from the same specification but using BA wages. Because the minimum wage is a small fraction of median BA wages in all regions, we do not expect to find any evidence of minimum wage effects. In this sense, this specification can be seen as a placebo test. The private sector sample passes this test: we do, in fact, find small estimated values of $\beta$ which are not statistically significant at any standard significance level. For the total sample (including the public sector), we obtain large negative, statistically significant estimated coefficients. This is the wrong sign for a spillover effect and arises because of the inclusion of public sector workers who are paid on a national scale. The complications arising from this are the main reason that the private sector sample is our preferred sample. Based on the results from that sample, minimum wages appear to have a spillover effect on HS median wages; an interpretation that is strengthened by the lack of estimated effect for the placebo, BA sample.

We are interested in how much of the relative movement in the HS wage outside of London we might explain through minimum wage spillovers. To capture this, we construct counterfactual HS wage series, $\tilde{y}_{rt}^{HS}$ for each region in which we hold minimum wage effects constant by subtracting predicted minimum wage effects from that actual wage in each year

Table 4: Graduate (BA) wage on the minimum-to-median bite (placebo)

| | All workers | | Private sector | |
|---|---|---|---|---|
| | All reg. | Excl. Ldn | All reg. | Excl. Ldn |
| $\text{NMW}_t/\text{Median}_{r,2010}$ | -1.791*** | -2.069** | 0.142 | 0.741 |
| | (0.270) | (0.891) | (0.469) | (0.682) |
| Observations | 192 | 176 | 192 | 176 |

Standard errors in parentheses

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

*Notes:* As Table 3 but for degree-holders, run on the graduate base bite. Region and year fixed effects; weighted by cell employment; SE clustered by region.

as follows:

$$\tilde{y}_{rt}^{HS} \;=\; y_{rt}^{HS} - \beta\big(b_{rt}^{HS,\text{act}} - b_{rt}^{HS,\text{cf}}\big) \;=\; y_{rt}^{HS} - \beta\,\frac{\text{NMW}_t - \text{NMW}_{2010}}{\text{Median}_{r,2010}^{HS}}, \qquad t \geq 2010, \tag{3}$$

where we use the estimate from column 3 of Table 3 (private sector, all regions, estimated on 2010-2025) for $\beta$. In Figure 12, we plot the non-London minus London HS wage differentials for both the actual and counterfactual series. The actual series shows the pattern we have seen before: by 2025 the non-London HS wage sits about 10% above its London counterpart relative to 2010, and essentially all of that gap opens up after 2014. The counterfactual series, which holds the real minimum wage at its 2010 value, is much flatter. By 2025 the actual gap stands at 0.103 log points against 0.014 in the counterfactual, so on these estimates the rising minimum wage accounts for roughly seven-eighths of the divergence between HS wages inside and outside London.

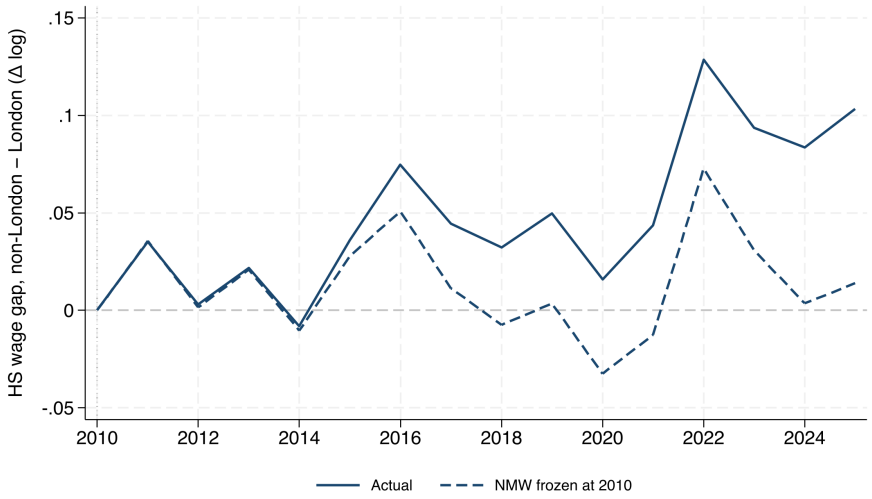


Figure 12: Non-London minus London HS wage gap, actual and minimum-wage-frozen counterfactual, private sector

*Notes:* LFS private-sector HS sample. The non-London−London gap in the age-adjusted log HS year effects: the solid line is the actual gap and the dashed line the gap with the real minimum wage frozen at its 2010 level. This is the difference of the two level lines in Figure 24.

In Appendix C, we present robustness checks. Most notably, we re-estimate our specification over alternative sample windows: the two disjoint sub-periods 2010-2015 and 2016-2025, which split the sample at the introduction of the National Living Wage, and the two pooled windows 2010-2025 and 2012-2025, which change the proportion of the sample accounted for by the strong post-2015 NMW rise. The base year of the bite variable is held at 2010 throughout. For the private sector sample, the four HS point estimates are all close to one another and to our main estimate (1.71, 1.64, 1.51 and 1.64). The three windows that contain the post-2015 minimum wage increases are each significant at the 1% level; the 2010-2015 sub-period is not distinguishable from zero, which is what we would expect given that the real minimum wage rose by only 3.3% over those years, leaving very little movement in the bite for the regression to use. The corresponding BA estimates are small and statistically insignificant in every window. We also repeat our estimation but work with the 75th percentile instead of the median. The estimated $\beta$ coefficient from this specification for private sector HS workers is 0.25 (se: 0.36), about a sixth of the corresponding median estimate and not distinguishable from zero. The lack of effect on the 75th percentile fits with our interpretation that we are capturing spillover effects since we would not expect spillovers to reach that high in the distribution. In contrast, if we were inadvertently capturing something like productivity trends then we would expect the 75th percentile to be affected as well.

As described earlier, the key identifying assumption is that there is no relationship between the level of the median wage in a base period and growth in the median wage in

subsequent years in the absence of minimum wage changes. We can obtain suggestive evidence for this by examining the relationship between the regional median HS wages in 2006 and growth in the median wages between 2006 and 2012. Because this is a period when the NMW is low relative to the regional median wages (it is over £3 below the median in most regions in this period), we argue that the patterns in this period reveal any underlying relationship between the level and growth in the median HS wage in the absence of minimum wage spillover effects. If we assume that the relationship revealed in these years is relevant for our sample period (2010-25) then this provides evidence about the identification assumption.

In Figure 13a, we plot the median wage growth from 2006 to 2012 against the median wage in 2006 for the private sector, HS sample. One can see in this graph how much of an outlier London is. Including London, there is a negative, statistically insignificant (at the 5% level) relationship (a slope of $-0.10$ with a standard error of 0.06). Excluding London, the relationship is almost flat with an estimated slope of $-0.03$ with an associated standard error of 0.13. We take this as suggestive evidence that our identification assumption is plausible, especially in the sample excluding London. Figure 13b repeats the exercise for 2013–2019, the period of the rising NMW, which starts just before the introduction of the National Living Wage and stops before COVID. In this period, lower-wage regions saw faster HS wage growth: the slope across all twelve regions is $-0.20$ (standard error 0.08), significant at the 5% level and about twice the 2006–2012 slope, and excluding London it is $-0.19$ (standard error 0.17), much steeper than in 2006–2012 although imprecisely estimated and not statistically significant. This contrast between the two windows is consistent with minimum-wage spillovers compressing regional HS wages in the later window.

### 7.2.2 Occupation x Region Variation

The evidence on spillovers presented so far suffers from relying on only 12 regional trends. To test the idea further, we turn to ASHE data in which there are enough observations to allow us to examine patterns by the interaction of occupation and region. More specifically, we work with the published mean and median hourly wages from ASHE that are reported in 4-digit occupation (SOC 2010) $\times$ region $\times$ year cell, yielding about ten thousand cells over 2012–2020 once we require a published median, a published employment count and a published 2012 base median in the cell.[20] Working with occupation data has the added benefit of providing a placebo: a minimum-wage spillover should appear in low-paid occupations, fade

[20]The ASHE data sample period is shorter than that for the LFS. The robustness exercise reported in the previous subsection indicates that altering the time period and the base on which we calculate our minimum wage bite variable does not alter our results when working with the LFS.

(a) 2006–2012

(b) 2013–2019

Figure 13: Regional HS median wage growth against the initial HS median, private sector, 2006–2012 and 2013–2019

*Notes:* LFS private-sector HS sample; region-level. The vertical axis is the log change in the real HS median wage over the window and the horizontal axis the log real HS median in the window's first year. The solid line is the OLS fit across all twelve regions and the dashed line the fit excluding London; both are weighted by first-year HS employment. Both panels share the same axis scales.

in middle-paid ones, and be absent for high-paid (BA-level) occupations that sit far above the floor.

To see whether the occupation x region data contains basic patterns that fit with the idea of spillover effects, in Figure 14a we plot, for each 4-digit occupation, its 2015–2019 real median wage growth against its 2012 real (base) median wage, coloured by the occupation's typical qualification level (NQF band). Each dot is a 4-digit occupation. The x-axis shows the occupation's median wage outside London, averaged across non-London regions in 2012. The y-axis shows the growth rate of that occupation's wage outside London, over 2015-19 in the left subfigure and over 2012-15 in the right one. The evidence supports the possibility of spillover effects from minimum wages. The lowest paid occupations show the most median wage growth over 2015-19, a period incorporating the substantial increase in the minimum wage in 2016. The minimum wage in 2012 was £6.11 in cash terms (the calendar-year average of the £6.08 and £6.19 rates), which is £8.88 in the 2025 prices used throughout. This is represented by a vertical line in the figure. The steepest wage growth after 2015 is in occupations whose 2012 median wage was within about £6 of that minimum – mainly low-skilled occupations – but the fitted line stays above zero out to a 2012 median of roughly £20, that is some £11 above the floor. The black line in the figure is a Lowess non-parametric line

and shows steep effects at low wages but quite a flat trend at higher wages. In occupations at the lowest end of the wage range, the real median wage grew by between 10 and 15 per cent in the four years between 2015 and 2019. For comparison, the increase in the real minimum wage over these same years was 14.8%. By contrast, real median wages in degree-level (NQF 6+) occupations fell slightly on average: across the 27 such occupations, growth over 2015–2019 averaged −0.016 log points (−0.009 when weighted by employment), and 70% of them recorded a fall.

To check further on whether the increases at lower wages in Figure 14a are plausibly related to minimum wage movements, in Figure 14b, we repeat the exercise but plot occupational wage growth over the 2012 to 2015 period against the median wage in the occupation in 2012. This is a period with only a 2.6% real increase in the minimum wage, and we see virtually no relationship between the base wage and wage growth as represented by the Lowess fit line. The contrast between the two figures fits with the post-2015 patterns arising due to minimum wage spillover effects.

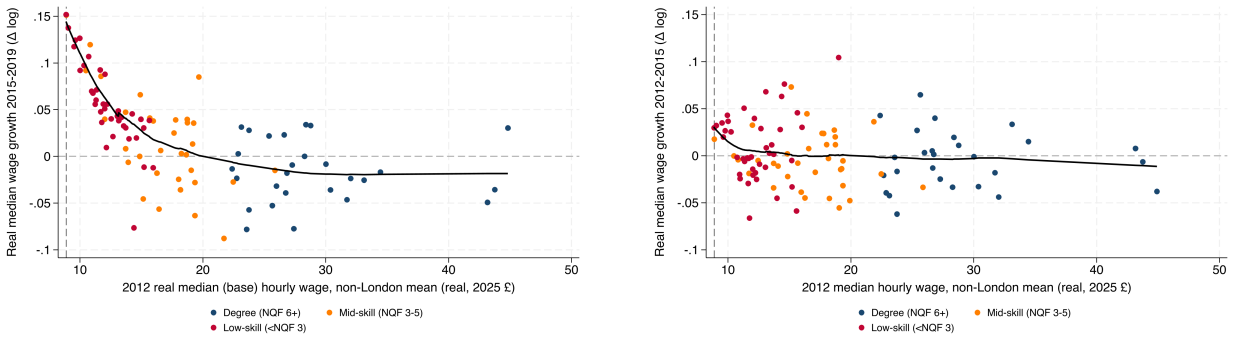


(a) 2015–2019 (b) 2012–2015

Figure 14: Occupation-level real median wage growth against the 2012 base median wage, by NQF band

*Notes:* ASHE, non-London Great Britain, in 2025 prices. Each point is a 4-digit SOC 2010 occupation observed in at least five of the ten non-London regions in 2012, 2015 and 2019 (98 occupations). The vertical axis is real median wage growth over 2015–2019 (left panel) and 2012–2015 (right panel); the horizontal axis is the occupation's 2012 real (base) median wage; points are coloured by NQF band. The dashed vertical line marks the 2012 real minimum wage (£8.88) and the black curves are lowess fits. The 2015–2019 window ends in 2019 to avoid the Covid-affected April 2020 ASHE.

We next estimate our panel wage regression using the occupation by region cells ASHE data for the years 2012 through 2020. In particular, we estimate:

$$\ln w_{jrt} = \eta_{jr} + \gamma_t + \beta \frac{\text{NMW}_t}{\text{Median}_{jr,2012}} + u_{jrt} \tag{4}$$

The inclusion of the complete set of occupation, j, by region, r, fixed effects ($\eta_{jr}$) implies

that the identifying variation is within occupation by region cells over time. As before, the assumption needed for identification is that the level of the median wage in the cell in 2012 is not correlated with future within cell wage growth except through the minimum wage spillover channel. We view Figure 14b as a check on that assumption, showing that in a period when the minimum wage wasn't increasing substantially, wage growth was very similar regardless of the initial median wage in the occupation. The only exception to this is a slightly higher growth rate in the very bottom occupations.

Table 5: ASHE occupation-level: log median wage on the bite, by NQF band

| | Degree | Mid-skill | Low-skill |
|---|---|---|---|
| $\text{NMW}_t/\text{Median}_{jr,2012}$ | 0.350 | 1.581*** | 2.079*** |
| | (0.680) | (0.099) | (0.115) |
| Observations | 2814 | 3490 | 3931 |

Standard errors in parentheses

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

*Notes:* Outcome is the log median hourly wage in 4-digit SOC 2010 × region × year cells, 2012–2020. Occupation × region and year fixed effects; employment-weighted; standard errors clustered by region.

Table 5 presents the results of the estimation, done separately by NQF band. The pattern across the different skill level occupations matches what we would predict from spillovers from the minimum wage increases that occur in the second part of our data period, namely: a small and statistically insignificant effect in the degree requiring occupations; a sizeable and statistically significant effect for mid-skilled occupations; and the largest effects for the lowest skilled occupations. Interestingly, the magnitude of the estimated effect for the middle skilled group is very similar to the effect we estimated earlier despite being based on an entirely different data source and using occupation x region rather than just regional variation. The large effects for the low-skilled occupations can help in understanding why the education wage premium series that uses all workers with an education below BA as the comparison group declines more than the other series that restrict the comparison group to be HS workers, however defined. Movements in the series using all below-BA workers will partly reflect the effects of the minimum wage on the wages of those without good GCSEs, whose wage levels are much closer to the NMW and, therefore, subject to more spillover effects.

### 7.2.3 Employment Effects

While we have argued that movements in the NMW could have had effects on HS wages in lower wage regions through fairness-related spillover effects, another possible channel is through selectivity of any lay-off effects. If minimum wage increases resulted in layoffs of the lowest paid workers in occupation x region cells where the minimum wage bite was largest then the median wage would move up in those cells, as observed. To check on this, in Figure 15, we recreate our growth versus base wage plot but examine employment rather than wage growth within occupations. The figure does show some negative employment growth in the lowest paid occupations, with small positive growth happening in occupations paying above about £13 in the base period, but the mean of the declines as represented by the non-parametric fit line is small and the changes are noisy.[21]

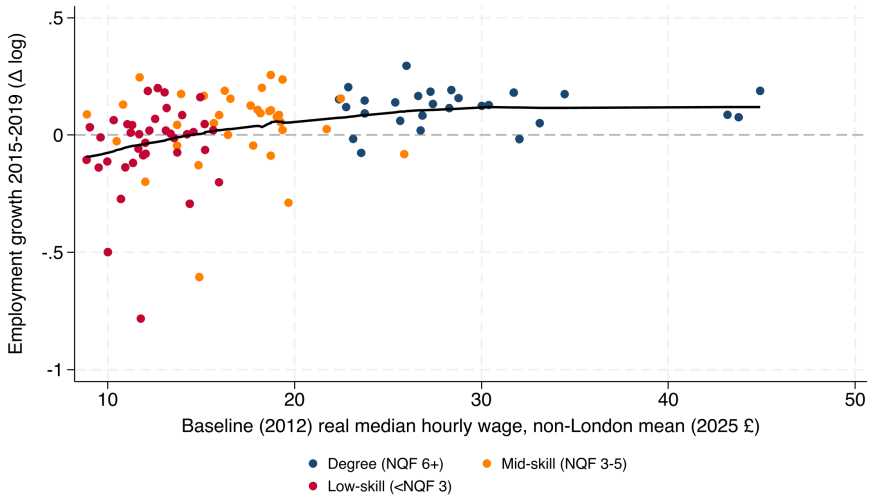


Figure 15: Occupation-level employment growth 2015–2019 against the 2012 base median wage, by NQF band

*Notes:* ASHE, non-London Great Britain. Each point is a 4-digit SOC 2010 occupation observed in at least five of the ten non-London regions. The vertical axis is log employment growth over 2015–2019 and the horizontal axis the baseline (2012) real median wage; points are coloured by NQF band and the black curve is a lowess fit.

To provide some more precision on the extent of any employment effects, we re-run our main specification replacing the log median wage in the cell as the dependent variable with log employment. We present the results, again run separately by skill level group, in Table 6. Within each group, the estimated effect is small and statistically insignificant.[22] The estimated coefficient is negative for the mid-skilled group, implying a small disemployment effect

[21] The three largest employment declines are shelf fillers (−0.78 log points, 2012 base median £11.8), national government administrative occupations (−0.61, £14.9) and retail cashiers and check-out operators (−0.50, £10.0); only the first two fall below −0.5.

[22] Because the regressor is the ratio $\text{NMW}_t/\text{Median}_{jr,2012}$, a coefficient can be put in more interpretable

from the minimum wage increases, but the coefficient for the lowest skill group, where we would expect to see the most layoffs, is actually positive (though very small and statistically insignificant). We view this evidence as fitting with the large literature finding negligible disemployment effects from minimum wage increases (Dube and Lindner, 2024). We also see it as implying that our wage results are not being driven by employment effects and, so, are more likely to stem from fairness-related wage spillovers.

Table 6: ASHE: log occupation employment on the bite, by NQF band

| | Degree | Mid-skill | Low-skill |
|---|---|---|---|
| $\text{NMW}_t/\text{Median}_{jr,2012}$ | 0.082 | -0.274 | 0.011 |
| | (0.946) | (0.389) | (0.243) |
| Observations | 2841 | 3504 | 3931 |

Standard errors in parentheses

$^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$

*Notes:* Outcome is the log of the occupation's 4-digit employment count in its region×year; the coefficient is a semi-elasticity (a one-unit rise in the bite maps to a 100×coefficient percent change in employment). Occupation × region and year fixed effects; unweighted; SE clustered by region.

Overall, our conclusion from the minimum wage exercises is that the greater rise in the HS wage outside versus in London after 2010 was plausibly driven by spillovers from the large post-2015 minimum wage increases. The minimum wage is national in application but it took values much closer to—and therefore much more likely to shift—HS median wages outside of London.

# 8 Investigating Other Forces

While minimum wage spillovers appear to provide a plausible explanation for the higher growth in the HS median wage outside of London, it seems possible that we could be inadvertently picking up the impacts of other forces. In this section, we examine the roles of a set of other possible driving forces and compare their impact on the HS wage patterns to the minimum wage impact. In particular, we consider differential movements (for regions outside London compared to London) in the shares of HS workers employed in each 1-digit

units by dividing it by the group's mean base median wage, which converts it into the effect of a £1 rise in the real minimum wage on log employment. For the mid-skill group the coefficient of −0.274, divided by a mean base median of £17.4, implies a fall in cell employment of about 1.6% for each £1 rise in the minimum wage; for the low-skill group the coefficient of 0.011, over a mean base median of £12.8, implies a negligible +0.1%. Both are small and statistically insignificant.

SIC92 industry in a region; the shares in female/private, male/public, and female/public employment cells; the shares in 1-digit SOC occupation groups; the share of immigrants; and the ratio of BA to HS workers. Apart from the BA/HS ratio, these variables capture potentially different compositional shifts over time. We will compare their effects to those of the minimum wage bite variable introduced in the previous section.

We carry out our investigation using an augmented version of the regression we used for the minimum wage investigation. In particular, we estimate:

$$y_{rt}^{h} \;=\; \alpha_r + \delta_t + \delta_t^{NL} + \gamma x_{rt} + \epsilon_{rt}, \tag{5}$$

where $y_{rt}^{h}$ is the log of the age-adjusted median wage for the HS education in region r in year t, $\alpha_r$ represents a complete set of region effects, $\delta_t$ is a complete set of time effects, $\delta_t^{NL}$ corresponds to the interaction of the year effects with a non-London region dummy (London is the omitted region and 2010 the omitted year, so each $\delta_t^{NL}$ is the non-London minus London gap in year $t$ relative to 2010), and $x_{rt}$ is a vector that includes values for the possible driving forces just mentioned plus the minimum wage bite variable. The observations are again the region×year cells for the 2010–2025 period. We are interested in the London specific time effects. If we do not include the $x$ vector, those effects reveal the differential movements in the HS wage in London versus the average of the other regions and are the same as the "actual" line in Figure 12. If the included elements of the $x$ vector fully explain the differential movement then the estimated London-time effects should move to zero.

It is well known that in an exercise such as this, the order in which we add in the elements of the $x$ vector can alter apparent conclusions about the importance of each factor. In response, we use the approach from Gelbach (2016) which provides an order-invariant decomposition. In the Gelbach approach, there is a primary covariate (or set of covariates) and a set of auxiliary controls. The question of interest is how the effect of the primary covariate changes as each of the auxiliary controls is introduced. Gelbach shows that the total effect of controlling for the auxiliary controls on the primary covariate coefficient can be captured through a standard omitted-variable bias formula. Because that "bias" term is additive in the contributing effects of the various auxiliary controls, one can assign proportions of the effect of controlling for the complete set of controls to each of the auxiliary covariates without any statement about ordering of introduction.[23] In the main example

[23]More formally, call the coefficient on the interaction of the non-London dummy and the 2025 year dummy when the $X$ matrix is not included, $\delta^{NL}_{2025,\text{base}}$ and the same coefficient when including the complete $X$ matrix, $\delta^{NL}_{2025,\text{full}}$. The decomposition allocates the difference across blocks of controls (e.g., the set of industry dummy variables) as $\delta^{NL}_{\text{base}} - \delta^{NL}_{\text{full}} = \sum_k \sum_{j \in k} \gamma_j \hat{\beta}_j$, where $\hat{\beta}_j$ is the coefficient on covariate $j$ in the full model and $\gamma_j$ is the non-London×2025 coefficient from an auxiliary regression of covariate $j$ on the base-model regressors. Every entry in the table is therefore the summed Gelbach component of a group of

in Gelbach (2016), the primary covariate is a black race indicator variable and the auxiliary covariates include education and a cognitive ability test score. The question is then how much the black worker wage disadvantage in US data is explained by education and cognitive ability and which of the auxiliary covariates is more important for altering the primary race effect. In our case, the primary covariate of interest is $\delta^{NL}_{2025}$, which captures the difference in the London/non-London HS wage differential between 2025 and the base period (2010). The question of interest is how much of the difference between London and the rest of the UK in HS wage movements over our sample period is reduced by controlling for industry composition, the bite of the minimum wage, etc.

Table 7 contains the results of our decomposition exercise. The left two columns present results associated with the complete set of controls minus the NMW bite variable. The right two columns add in the bite variable. The top row gives the estimated $\delta^{NL}_{2025,\text{base}}$ coefficient, which the table writes as $\theta_{\text{base}}$ (with $\theta_{\text{full}}$ for $\delta^{NL}_{2025,\text{full}}$), and indicates that non-London HS wages grew, on average, 0.103 log points faster than London HS wages between 2010 and 2025. The "Total explained" row in the first column indicates that the set of covariates other than the bite variable can explain .028 log points out of the total, or 27.5% of the total difference in the movements of HS wages out versus in London. The intermediate rows show the contributions of the various auxiliary control variable blocks. The gender, public-sector and gender × public-sector blocks together contribute 0.015 log points, and each is statistically significant at the 5% level; the industry (0.006) and occupation (0.008) blocks are of similar size but are not statistically distinguishable from zero.[24] Note that, matching our earlier exercise, differences in movements in the BA/HS supply explain very little of the relative wage movements.

The right two columns repeat the decomposition exercise but add in the minimum wage bite variable. The bite variable on its own over-explains the relative rise in HS wages outside London, implying that HS wages outside of London would actually have fallen if the NMW had remained at its 2010 level. This contrasts with our earlier minimum wage exercise in which movements in the NMW almost exactly account for the relative rise in non-London HS wages. As seen at the bottom of the table, the estimated bite variable effect is larger in the current regression, indicating that the coefficient in the earlier exercise was biased by the omission of the industry, etc. variables. With the bite variable included, the Gender/public composition block accounts for almost the entire non-minimum wage related impact while the occupation and industry blocks have off-setting effects and the remain-

variables rather than a regression coefficient. The unexplained row is the non-London minus London gap that survives the controls.

[24] Examining the individual industry effects underlying the total industry block indicates that the industry contribution is not a low-paid-services story: it is led in roughly equal measure by *financial intermediation* and *wholesale and retail trade*, whose employment shares both rose outside London relative to within it.

Table 7: Gelbach decomposition of the 2010-2025 change in the non-London−London HS wage gap, private sector.

| | Composition only | | With NMW bite | |
|---|---|---|---|---|
| | (1) | | (2) | |
| | Log pts | % of gap | Log pts | % of gap |
| Baseline gap $\theta_{\text{base}}$ | 0.103 | 100.0 | 0.103 | 100.0 |
| Industry | 0.006 | 5.7 | -0.007 | -6.5 |
| | (0.012) | | (0.013) | |
| Gender (female share) | 0.009 | 8.6 | 0.009 | 9.0 |
| | (0.003) | | (0.002) | |
| Sector (public share) | 0.009 | 8.2 | 0.008 | 7.5 |
| | (0.004) | | (0.004) | |
| Gender × sector | -0.003 | -3.2 | -0.003 | -2.8 |
| | (0.001) | | (0.001) | |
| Occupation | 0.008 | 7.8 | 0.012 | 11.4 |
| | (0.007) | | (0.007) | |
| Immigration | 0.001 | 0.9 | -0.001 | -1.1 |
| | (0.001) | | (0.000) | |
| BA/HS supply | -0.001 | -0.5 | -0.000 | -0.5 |
| | (0.000) | | (0.000) | |
| NMW bite | — | — | 0.123 | 118.9 |
| | | | (0.018) | |
| Total explained $\theta_{\text{base}} - \theta_{\text{full}}$ | 0.028 | 27.5 | 0.141 | 135.8 |
| | (0.011) | | (0.014) | |
| Unexplained $\theta_{\text{full}}$ | 0.075 | 72.5 | -0.037 | -35.8 |
| | (0.023) | | (0.020) | |
| *Memo: NMW-bite slope in column (2)* | | | | |
| NMW bite coefficient $\hat{\beta}$ | — | | 2.074* | |
| | | | (0.403) | |
| wild-bootstrap $p$ | — | | 0.056 | |
| wild-bootstrap 95% CI | — | | [-0.078, 4.105] | |

*Notes:* Gelbach (2016) decomposition of the non-London×2025 interaction coefficient, in log points and as a share of the baseline gap; the identity and the construction of each entry are given in the text. Employment-weighted, 2010–2025. Standard errors in parentheses: wild-cluster bootstrap (Webb weights, 9999 replications) for the blocks and totals, clustered by the twelve workplace regions for $\delta^{NL}_{2025,\text{full}}$ and the bite coefficient.

ing factors have negligible effects. The bite component is much the largest and is precisely estimated (0.123, standard error 0.018). Among the composition blocks, the female-share (0.009) and gender × public-sector (−0.003) components are clearly distinguishable from zero and the public-sector component (0.008, standard error 0.004) is borderline, but all of them are an order of magnitude smaller than the bite component. Including the bite variable reduces the contributing effects of the other covariates substantially, implying some level of collinearity between the minimum wage variable and the other variables. Nonetheless, the coefficient on the bite variable reported at the bottom of the table remains large and is significant at the 10% level (wild-bootstrap $p = 0.056$), suggesting that there is still enough independent variation to identify the minimum wage effect separately.

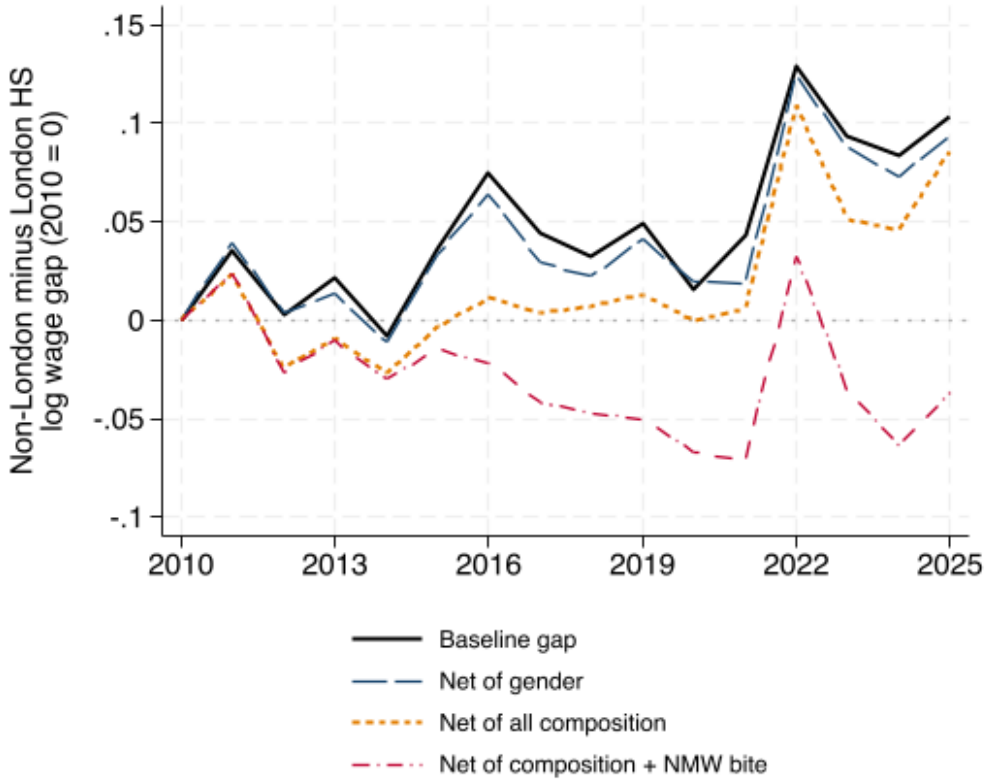


Figure 16: Cumulative decomposition of the non-London HS wage gap: composition versus the minimum-wage bite

*Notes:* LFS private-sector HS sample, employment-weighted, region and year fixed effects plus non-London×year interactions, over 2010–2025; London = control and 2010 = 0. Each line is the baseline London/non-London gap net of the cumulative Gelbach contributions of the labelled blocks, taken from the full specification with the minimum-wage bite (column (2) of Table 7): the solid line is the baseline gap; the next removes the female-share block; the next removes all composition controls; the last also removes the minimum-wage bite and coincides with the fitted full-model gap. Bite = $\text{NMW}_t/\text{Median}^{HS}_{r,2010}$. The plotted paths are point estimates; the underlying regressions are estimated with standard errors clustered by region.

Figure 16 traces this decomposition through time, cumulatively peeling the estimated contributions off the baseline gap using the full specification with the minimum-wage bite. The top line in the figure shows the baseline gap in each year (i.e., the gap between Non-London and London HS wages when no extra controls are included). We construct the other lines by sequentially including the parts of the Gelbach decomposition associated with named blocks of auxiliary controls. Thus, the "Net of gender" line adds the decomposition component associated with the female-share block to the $\delta^{NL}_{t,\text{base}}$ coefficients but leaves the other decomposition elements out, and the "Net of all composition" line adds the components for the whole set of composition controls. The plot indicates that gender composition on its own accounts for very little of the differential HS wage movements outside versus in London, but that the composition controls taken together account for most of those movements before 2020 and only a small part of the subsequent further rise in non-London wages after 2020. The minimum wage movements can also account for the relative increase in non-London HS wages before 2020 on its own as well as the further increases after 2020. Together, the compositional shifts and the minimum wage spillovers over-explain the relative wage differential across our sample period. It is worth noting that if we focus just on the post-2014 period (the period over which the HS wage outside of London actually increases relative to the London HS wage), the minimum wage variable only slightly over-explains the actual wage differential movement (accounting for 107% of the movement) and the other components as a group explain none of the movement. Thus, the other components over-explain movements between 2010 and 2014 (a period with little movement in the minimum wage) but have no additional effect thereafter. It is the minimum wage that is driving trends in the period when HS wages inside and outside of London truly diverge.[25]

We view this exercise as showing that the relative rise in HS wages outside versus inside London is accounted for statistically mainly by the minimum wage bite, with some contribution from shifts in the public sector x gender composition of the workforce, though some caution is warranted in that conclusion given the wide wild-bootstrap confidence sets. We view both of these components as quite different from what has been emphasised in standard explanations for the declining BA/HS wage premium - stories related to shifts in relative supply and industry or occupation composition.

# 9 Conclusion

Movements in the wage premium for a university education in the UK have attracted a lot of attention in recent years. Evidence that the premium has been declining has been argued to

[25] We present the decomposition results using a 2014 normalisation in Appendix D.

be directly linked to productivity problems, with the claim that creating and adapting new technologies should be creating new demand that would push the premium up. Declines, instead, seem to imply lagging technological adaptation that has meant that demand for university workers has not kept up with the surging supply created by the expansion of the university system. The fact that the decline in the premium has happened predominantly outside London (as Xu (2023) and Stansbury et al. (2023) document and we find also) seems to support this outlook since productivity problems have been particularly acute in the regions of the UK outside of London.

In this paper, we draw on LFS and ASHE data to provide a careful examination of movements in the UK university wage premium over the last three decades. We begin by establishing that, even in the face of a very large expansion in the number of university graduates, the premium was unchanged between 1997 and about 2010 but has been declining since 2010. We show that this pattern mostly stands up to working (as the literature does) with different definitions of the comparison group. Typically, the wage for university graduates (which we call BA's) is compared to that for high school graduates (HS's) and different definitions of who is a high school graduate do not alter this overall pattern. If, however, we include the bottom education group, those without good GCSEs, in the comparison group then the premium has been declining since 1997. We argue that in order to discuss the implications of the wage premium for productivity, as the literature does, we need a well-defined comparison group. That and the likelihood that the wages for the very lowest education group are being pushed around by the minimum wage lead us to focus on the BA - HS wage premium with its flat then declining pattern.

As both Xu (2023) and Stansbury et al. (2023) highlight, the movements in the premium have an important regional dimension. The wage premium stays constant throughout the 1997-2025 period in London but for the other regions of the UK it is flat until 2010 then declines afterward. We show that the relative decline after 2010 for non-London regions versus London comes mainly from an unexpected source: it is largely driven by a rise in the HS wage outside of London relative to in London. Once wages are age-adjusted, the BA wage outside of London shows no significant decline relative to London after versus before 2010. As Xu (2023) suggests, BA workers outside of London have become more likely than those in London to work at lower skilled jobs since 2001, and graduate wages are strongly related to such occupational placement. But while the tendency of university workers to work at lower skilled jobs is much higher outside than in London at all times in our sample period, the relative change in that difference has been small since 2010, and it does not show up as a relative decline in non-London graduate wages. In other words, occupational downgrading is real, but it cannot account for the decline in the premium. Thus, understanding what is driving the differences in the movements in the HS wage inside versus outside London is the

key to understanding movements in the BA/HS wage premium.

We argue that the most plausible explanation for the rise in the HS wage outside of London is the rise in the National Minimum Wage (NMW). In 2010, the NMW was at least £3 below the median private-sector HS wage in 9 regions in the UK, but with the increase in 2016, it moved to within £2 to £3 of the median in multiple regions and even within £2 for some. The idea that the minimum wage could have effects on the median HS wage when it gets this close to that wage fits with results in the small literature examining wage spillover effects from minimum wages (Cengiz et al., 2019; Gopalan et al., 2021; Giupponi et al., 2024; Brochu et al., 2025). We show that such spillovers are a plausible explanation for movements in the HS median wage outside of London using both regional and occupation x region variation and working with both LFS and ASHE data. On the other hand, the fact that the median HS wage in London is always at least £5 above the NMW can explain why that wage doesn't rise in the same way as the non-London HS wage. Interestingly, the pushing up of median wages outside of London by the NMW increases was not accompanied by significant employment declines in the low- and mid-skilled occupations where the wage effects are concentrated.

It is worth emphasising, at this point, a feature of the movements in the BA/HS premium that is most at odds with expectations: the lack of movement in the premium across the whole 1997–2025 period in London and over the period 1997 to 2010 outside of London. This is unexpected because of the very substantial rise in the proportion of the workforce with a BA over this same time period. In a standard demand and supply framework, one would expect the large increase in the relative supply of BA workers to have driven down the BA-HS premium. Indeed, this is the explanation typically offered for the decline in the premium after 2010. We show that that explanation doesn't work for the UK. The regions where the relative supply grew fastest are not the ones where the BA/HS premium declined the most. The regions where it declined the most are the ones with the lowest initial HS wages, which is where minimum wage spillovers would be expected to be strongest.

The minimum wage bite accounts statistically for the relative rise in HS wages outside London, but we cannot rule out all other explanations. Our analysis cannot exclude explanations beyond those captured by our control variables (industry, occupation, gender and public-sector composition, the immigrant share, and the BA/HS ratio). For example, the supply of lower- and middle-skilled labour may have shifted differently across regions after 2016, due to reduced EU migration, and our immigrant-share control may capture the regional incidence of such a shift only partly. Any other force after 2015 that raised low- and middle-level wages most in the regions where the 2010 HS median was lowest would load onto our bite variable, because that is exactly the pattern on which our identification relies.

We regard minimum wage spillovers as the leading candidate, but other possibilities remain open.

So, we are left with a somewhat different picture from that often put forth in many policy discussions. The university wage premium did not decline at all up until 2010 in the UK and remained flat in London after 2010. If minimum wage spillovers are indeed the main force behind the rise in HS wages outside London, the premium in the other regions might also have remained flat in the absence of the substantial increase in the minimum wage. In Blundell et al. (2022), we provide evidence in support of an explanation for why the substantial increase in the supply of BA's could have been accompanied by no change in the BA/HS wage premium. We argue that such a pattern can arise in situations in which firms are choosing what technology to adopt and are induced to select more BA intensive technologies by the increase in the relative supply of BA's. In an echo of a standard trade model, the increased supply of BA workers essentially creates its own, offsetting demand. This seems like a less dire depiction of technological implementation than is sometimes discussed. However, it may fit best with UK firms taking up technologies already created elsewhere, as opposed to being innovators.

In the end, our results do not calm concerns about productivity issues in the UK, including their regional imbalances. But they do suggest that the recent decline in the BA/HS premium is not a bellwether for productivity failures. Instead, it most plausibly reflects the impact of minimum wage policy, and it raises an alternative set of questions: what would have happened to the premium outside of London in the absence of the minimum wage increases, and how have the minimum wage increases affected technology adoption and creation by firms outside of London? We leave investigating those questions for future work and debate.

# References

Barbas, N., A. Valero, and J. Van Reenen (2026, July). UK national productivity trends: Getting the data right. CEP Occasional Paper 74, Centre for Economic Performance, London School of Economics, London.

Blundell, R., D. A. Green, and W. Jin (2022). The u.k. as a technological follower: Higher education expansion and the college wage premium. *Review of Economic Studies 89*(1), 142–180.

Britton, J., K. Ogden, N. Ornadel, and B. Waltmann (2026, June). New estimates of the impact of undergraduate degrees on lifetime earnings. Ifs report, Institute for Fiscal Studies, London. ISBN 978-1-83870-818-4.

Brochu, P., D. A. Green, T. Lemieux, and J. Townsend (2025). The minimum wage, turnover, and the shape of the wage distribution. *Journal of Labor Economics*. forthcoming.

Burn-Murdoch, J. (2026, August). To fix education, fix the economy first. *Financial Times*.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics 120*, 1405–1454.

Dube, A. and A. Lindner (2024). Minimum wages in the 21st century. In C. Dustmann and T. Lemieux (Eds.), *Handbook of Labor Economics*, Volume 5, pp. 261 – 383. Elsevier.

Gelbach, J. B. (2016). When do covariates matter? and which ones, and how much". *Journal of Labor Economics 34*(2), 509–543.

Giupponi, G., R. Joyce, A. Lindner, T. Waters, T. Wernham, and X. Xu (2024, April). The employment and distributional impacts of nationwide minimum wage changes. *Journal of Labor Economics 42*(S1), S293–S333.

Gopalan, R., B. Hamilton, A. Kalda, and D. Sovich (2021). State minimum wages, employment, and wage spillovers: Evidence from administrative payroll data. *Journal of Labor Economics 39*(3), 673–707.

Home Office, U. V. and Immigration (2014, April). Codes of practice for skilled workers: Standard occupational classification (SOC) codes. Version 04/14, UK Visas & Immigration, Home Office, London. Archived document; 6 April 2014.

Office for National Statistics (2025a). Annual Survey of Hours and Earnings (ASHE). Office for National Statistics. Crown Copyright. Employer-reported earnings and hours from an approximately 1% sample of employee jobs drawn from HMRC PAYE records.

Office for National Statistics (2025b). Quarterly Labour Force Survey. [data series]. UK Data Service [distributor]. Crown Copyright. Accessed under the UK Data Service End User Licence Agreement. Individual quarterly study numbers (SN) and DOIs are available from the UK Data Service, `https://ukdataservice.ac.uk`.

Office for National Statistics (2026a). Component approach to measuring labour productivity, UK: 1997 to 2024. ONS methodology article. `https://www.ons.gov.uk/economy/economicoutputandproductivity/productivitymeasures/methodologies/componentapproachtomeasuringlabourproductivityuk1997to2024`.

Office for National Statistics (2026b). Investigating non-response to the Labour Force Survey using administrative data. ONS methodology article, 28

September 2026. https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/methodologies/investigatingnonresponsetothelabourforcesurveyusingadministrativedata.

Office for National Statistics (2026c). Labour force survey performance and quality monitoring report: January to march 2026. released 19 May 2026, ONS website.

Stansbury, A., D. Turner, and E. Balls (2023). Tackling the UK's regional economic inequality: binding constraints and avenues for policy intervention. *Contemporary Social Science 18*(3-4), 318–356.

Xu, X. (2023, November). The changing geography of jobs. Technical Report R286, Institute for Fiscal Studies, London.

Xu, X. (2025). On the move: Local opportunities and youth migration. *Fiscal Studies 46*, 451–473.

# A Additional descriptive figures

This appendix supports nine claims made in the main text: (i) that our headline BA–HS wage-differential patterns are not sensitive to how we adjust for the changing age composition of the workforce; (ii) that they are not driven by changes in the gender composition of the education groups either, nor is the minimum-wage estimate of Table 3; (iii) that the flatness of the BA differential in London holds across all of the alternative high-school definitions considered in Figure 4; (iv) that the BA–HS differential of Figure 4 can be reconciled with the graduate premium series reported by Stansbury et al. (2023); (v) that the post-2010 decline in the non-London differential of Figure 3 is much smaller in the public sector; (vi) that the education-group wage patterns of Figure 5 are a private-sector phenomenon; (vii) that the occupation-based wage patterns we read off the ASHE in Figure 8 are also present in the LFS; (viii) what the minimum-wage counterfactual of equation (3) looks like in levels rather than as a London/non-London difference; and (ix) what the gap bands of Figure 10 look like when the gap is measured in constant prices, and what they look like in the public sector.

## A.1 Age adjustment: age-cell means versus a Mincerian regression

Throughout the paper we age-adjust the BA–HS median wage ratio by collapsing to (year × five-year age band × education) cells and regressing the cell-level log median ratio on age-

band and year dummies (Section 2.2). Stansbury et al. (2023) instead age-adjust with an individual-level Mincerian log-wage regression. Figure 17 shows that this choice is immaterial for the UK-wide private-sector series: our baseline age-cell median-ratio series, the individual-level Mincerian BA coefficient, and an age-cell mean-log-wage-gap series all trace the same path: flat through about 2010 and declining thereafter.

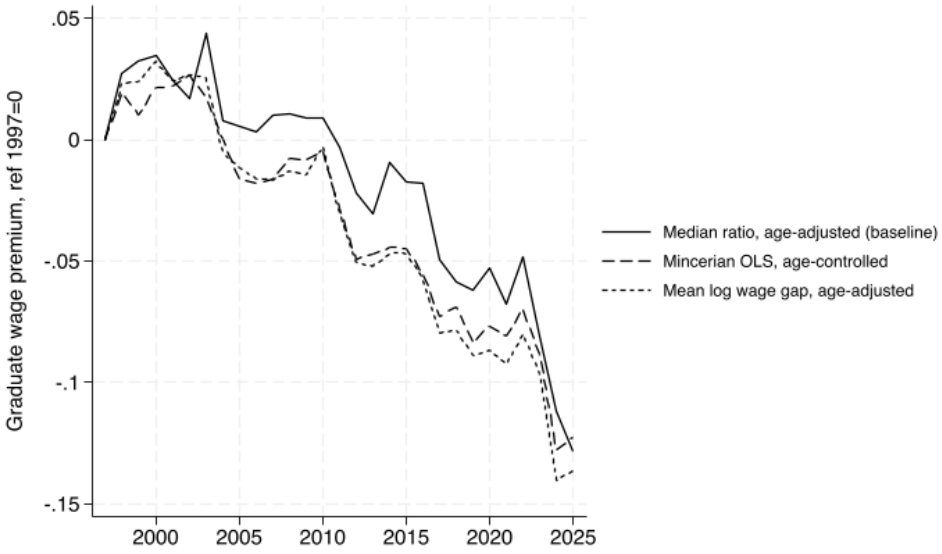


Figure 17: The UK private-sector age-adjusted BA/HS wage differential under three age-adjustment methods, 1997–2025

*Notes:* LFS private-sector employees aged 20–59, 1997–2025, actual-hours hourly wage, reference 1997 = 0. The series starts in 1997 because the LFS sector variable is populated for only about a quarter of the 1993 wage records, too few to anchor a private-sector series. The three series differ only in how they adjust for age composition: the age-band cell log median-ratio series used throughout the paper (baseline); an individual-level Mincerian regression of log real hourly wage on a BA dummy and five-year age-band dummies, estimated year by year (the coefficient on the BA dummy); and an age-band cell mean-log-wage-gap series. HS = GCSE grade C or above without a degree; BA is degree or higher.

Figure 18 repeats the Mincerian regression separately for London and the pooled rest of the UK (by region of workplace), on the private-sector sample used for the paper's main results. The regional divergence that is central to the paper survives intact: the graduate premium is essentially flat in London across the whole period, whereas outside London it declines steadily after 2010. Because this holds under the Mincerian age adjustment combined with our GCSE C+ definition of the high-school group, the difference between our series and that reported by Stansbury et al. (2023) does not stem from our method of age adjustment.

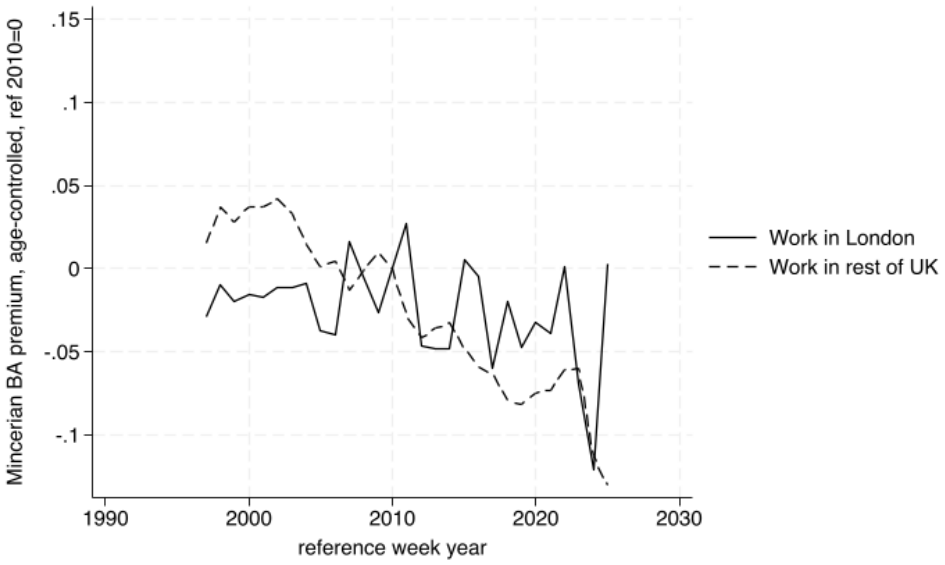


Figure 18: Mincerian BA/HS wage differential, London versus rest of the UK, 1997–2025

*Notes:* LFS private-sector employees aged 20–59, 1997–2025, by region of workplace, actual-hours hourly wage, reference 2010 = 0. Each line is the year-by-year Mincerian BA coefficient (log real hourly wage regressed on a BA dummy and five-year age-band dummies), estimated separately for London and the pooled rest of the UK. The workplace region is available from 1997. HS = GCSE grade C or above without a degree; BA is degree or higher.

## A.2 Gender composition

Figures 2 and 3 adjust for age but pool men and women within each age cell, so a change in the sex mix within an education group moves the series. Among non-London private-sector wage earners, the female share among graduates rose from 39% in 2010 to 47% in 2025, while among HS workers it fell slightly, from 44% to 42%. We therefore re-estimate the private-sector series on (year × five-year age band × sex × education) cells, regressing the cell log BA/HS median ratio on a full set of age-band × sex dummies and year dummies. Figure 19 shows the result. Outside London the age-sex-adjusted differential is −0.077 in 2019 and −0.135 in 2025, against −0.100 and −0.155 in Figure 3, so gender composition accounts for about 2 of the 15.5 log point decline. London stays flat (−0.026 in 2025, against −0.031). Because the cells are weighted equally, the adjusted series is close to the average of the male and female premia. Splitting by sex halves the cells: the London private-sector cells have a median of about 65 wage observations for BA and 69 for HS, with some cells below 10, so the London series is noisier; the non-London cells have medians of 276 (BA) and 765 (HS), and none has fewer than 20.

The dependent variable of Table 3 is likewise age- but not sex-adjusted. Table 8 re-estimates it with the age-sex-adjusted HS wage. The private-sector coefficients are 1.450

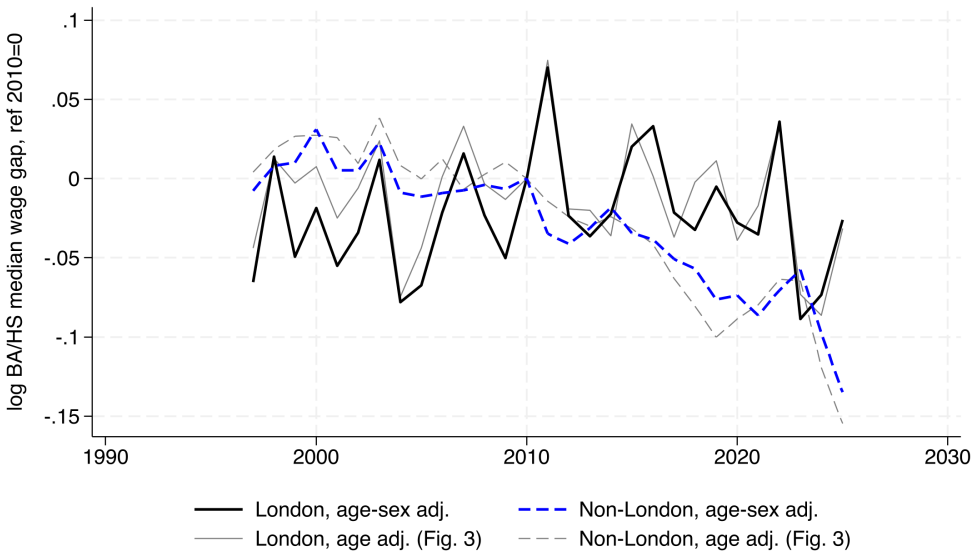


Figure 19: Age-sex-adjusted log BA/HS median wage ratio, private sector, London versus rest of the UK, 1997–2025

*Notes:* LFS private-sector employees aged 20–59, 1997–2025, by region of workplace; actual-hours hourly wage; reference 2010 = 0. Thick lines (black solid: London; blue dashed: rest of the UK): year effects from an unweighted regression of the log BA/HS median wage ratio in each year × five-year age band × sex cell on age-band × sex dummies and year dummies. Thin grey lines: the age-adjusted series of Figure 3 (solid London, dashed rest of the UK).

(s.e. 0.119) for all regions and 1.518 (0.219) excluding London, against 1.513 and 1.600 in Table 3, so the column-3 coefficient used in the counterfactual of equation (3) falls by about 4%. The all-worker, all-region coefficient falls more, from 1.202 to 0.955, but remains significant at the 1% level; excluding London it is 0.575 against 0.620, insignificant in both. Region × age band × sex cells are thin in the private sector (a median of 44 HS wage observations; 239 of 3,455 cells below 10), but cells are weighted by size in the first-stage regressions.

Table 8: HS wage on the minimum-to-median bite: age-sex-adjusted HS wages

| | All workers | | Private sector | |
|---|---|---|---|---|
| | All reg. | Excl. Ldn | All reg. | Excl. Ldn |
| $\text{NMW}_t/\text{Median}_{r,2010}$ | 0.955*** | 0.575 | 1.450*** | 1.518*** |
| | (0.166) | (0.374) | (0.119) | (0.219) |
| Observations | 192 | 176 | 192 | 176 |

Standard errors in parentheses

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

*Notes:* As Table 3 except the outcome: the age-sex-adjusted year effect of the median HS log real wage, from region-specific regressions of region × five-year age band × sex × year cell medians on age-band × sex dummies and year dummies, weighted by cell size (summed survey weights). Same bite $\text{NMW}_t/\text{Median}_{r,2010}$, region and year fixed effects, weighted by cell employment, 2010–2025, ages 20–59; standard errors clustered by region (12 clusters in the all-region columns, 11 in the columns excluding London). Produced by `syntax/NMW and wages by EDU region v4.do`, Section 3e.

## A.3 Alternative high-school definitions: London

Figure 4 in the main text shows the private-sector BA–HS differential outside London under three alternative definitions of the high-school comparison group. Figure 20 is the London counterpart, on the same private-sector sample. In contrast to the non-London picture, the London differential is essentially flat across the entire period under every one of the three definitions—including the broadest "all non-BAs" definition, which incorporates workers below GCSE grade C. The series oscillate rather than trend; the one qualification is that the A-levels-only line sits below the other two through much of the period, most visibly over 2003–2006 and again from 2013, and ends 2025 about 0.08 log points below its 2010 value against 0.03 for our preferred definition. This reinforces the conclusion that the post-2010 decline in the BA–HS differential is a non-London phenomenon and is not an artefact of how the high-school group is defined.

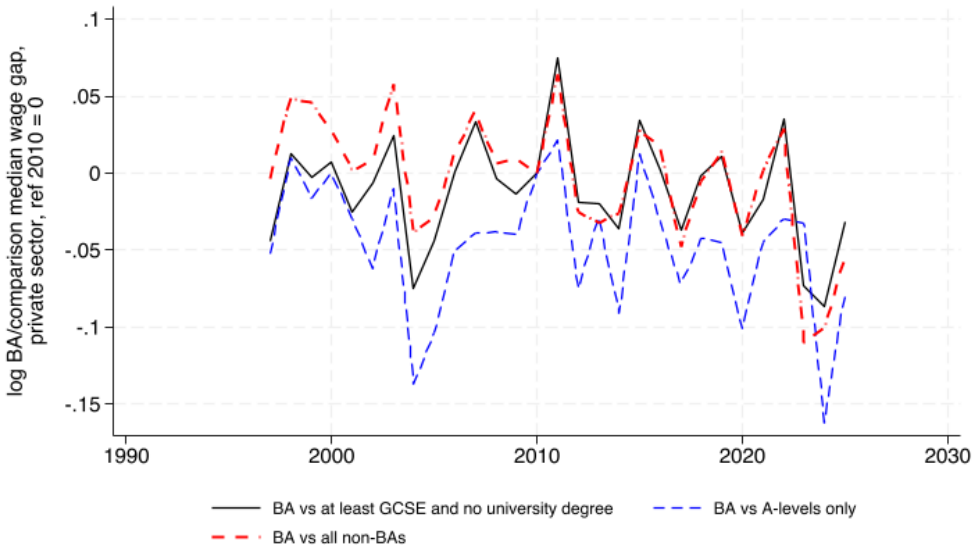


Figure 20: Age-adjusted BA–HS wage differential in London for private-sector workers under alternative high-school definitions, 1997–2025

*Notes:* LFS private-sector employees, London (by region of workplace), persons aged 20–59, actual-hours hourly wage. Each line is the age-adjusted log ratio of the BA median wage to the comparison-group median wage (log ratio in each year × five-year age band cell regressed on age-band and year dummies), reference 2010 = 0. Each line uses a different high-school definition: (1) GCSE grade C or above without a degree, including A-levels and sub-degree tertiary (Blundell et al. 2022); (2) A-levels only (Stansbury et al. 2023); and (3) any qualification below a degree, including those below GCSE grade C. The BA group is degree or higher throughout. This is the London counterpart of Figure 4.

## A.4 Reconciliation with Stansbury et al. (2023)

Figure 4 shows the BA–HS differential outside London as flat between 1997 and 2010. Stansbury et al. (2023) plot a private-sector graduate premium for each region that appears to fall across the whole of their 1997–2019 window. This subsection reconciles the two. Their construction differs from Figure 4's in six respects, and we vary them one at a time on our data. Table 9 reports each rung; the column to read is the 1997–2010 change, which is unaffected by how a series is normalised.

The line of Figure 4 that uses our preferred GCSE C+ comparison group moves by −0.004 log points between 1997 and 2010. The series plotted in Stansbury et al. (2023)'s Figure 5, rebuilt on our data, moves by −0.043. The 0.039 difference decomposes as follows.

- *The comparison group* (row 2). Replacing GCSE C+ with A-levels only—the omitted category in their regression—accounts for −0.020. This is the A-levels line of Figure 4. Stansbury et al. (2023) report that their result is robust to using GCSEs instead, so they do not regard this as the source of the difference.

- *The age floor* (row 3). They use 25–59 year olds where we use 20–59. This works in the opposite direction, accounting for +0.019, and so almost exactly cancels the comparison-group effect. Measured inside their own regression specification rather than ours the offset is smaller, +0.007 (rows 6d and 6).

- *The estimator* (row 4). They estimate the premium as the coefficient on a degree dummy in an individual-level log-wage regression; we take the ratio of group medians within age cells. This accounts for −0.020. Both approaches control for age, so what matters here is the difference between a mean of logs and a ratio of medians, not the method of age adjustment.

- *The sex control* (row 5). Adding a sex dummy accounts for a further −0.015.

- *Regional aggregation* (row 6). Estimating the premium region by region and averaging the eleven coefficients, rather than pooling the non-London regions, accounts for +0.020, offsetting the sex control.

- *Smoothing* (row 7). Their figure plots a five-year centred moving average of the premium. This accounts for −0.023, the largest single term.

Four of the six steps therefore cancel in two pairs: the comparison group against the age floor, and the sex control against the regional aggregation. What survives is the estimator and the smoothing, which between them account for −0.043—the whole of the gap.

Table 9: Reconciliation ladder: the non-London BA versus A-levels-only wage differential under Stansbury et al. (2023)'s construction

| Series | 1997 | 2010 | 2019 | Chg. 97–10 | Chg. 10–19 |
|---|---|---|---|---|---|
| *Panel A. Sequential ladder, private sector, non-London: one change per row* | | | | | |
| (1) BA vs GCSE C+ no degree, ages 20–59, cell-median ratio (2010 = 0) | 0.004 | 0.000 | -0.100 | -0.004 | -0.100 |
| (2) as (1), comparison group A-levels only (2010 = 0) | 0.024 | 0.000 | -0.089 | -0.024 | -0.089 |
| (3) as (2), ages 25–59 (2010 = 0) | 0.005 | 0.000 | -0.107 | -0.005 | -0.107 |
| (4) as (3), mean-log OLS in place of cell-median ratio (premium level) | 0.396 | 0.372 | 0.278 | -0.025 | -0.094 |
| (5) as (4), plus a sex control (premium level) | 0.419 | 0.379 | 0.291 | -0.040 | -0.088 |
| (6) as (5), region by region in place of pooled non-London (premium level) | 0.389 | 0.369 | 0.286 | -0.020 | -0.083 |
| (7) as (6), 5-year centred moving average (premium level) | 0.390 | 0.347 | 0.295 | -0.043 | -0.052 |
| *Panel B. Deviations from row (6), one at a time; NOT steps of the ladder* | | | | | |
| (6a) as (6), all sectors in place of private only (premium level) | 0.402 | 0.395 | 0.302 | -0.008 | -0.092 |
| (6b) as (6), usual-hours wage (premium level) | 0.355 | 0.369 | 0.282 | 0.015 | -0.087 |
| (6c) as (6), no sex control (premium level) | 0.368 | 0.360 | 0.271 | -0.008 | -0.089 |
| (6d) as (6), ages 20–59 (premium level) | 0.357 | 0.331 | 0.264 | -0.027 | -0.067 |

*Notes:* LFS, non-London regions by region of workplace, private-sector employees, actual-hours hourly wage, 1997–2025. Panel A is a sequential ladder: each row changes exactly one element of the row above, starting from the line of Figure 4 that uses our preferred GCSE C+ comparison group and ending at the series Stansbury et al. (2023) plot in their Figure 5. Rows (1)–(3) are indices normalised to zero in 2010 and rows (4)–(7) are premium levels, as each label states; the two change columns are invariant to that normalisation and are the entries that are comparable across rows. Rows (1) and (2) reproduce the corresponding lines of Figure 4 exactly. Row (6) reproduces Stansbury et al. (2023)'s specification: the log hourly wage regressed year by year on a degree dummy interacted with region of workplace, education dummies with A-levels omitted, and age-band and sex dummies, with the non-London premium taken as the mean of the eleven regional coefficients. Because that specification estimates region by region, rows (6) and (7) are fitted on the full UK sample including London, whereas rows (1)–(5) are fitted on the non-London regions alone. Row (7) applies the five-year centred moving average that their Figure 5 plots. Panel B varies one element of row (6) at a time and its rows are not steps of the ladder; the usual-hours row also changes the estimation sample, because the usual-hours and actual-hours wages are non-missing for different observations.

The smoothing effect is mechanical. A five-year centred moving average at year $t$ averages the years $t-2$ through $t+2$, so the 2010 point of the plotted series is the mean of the 2008–2012 values, two of which lie after the break. Because the premium falls steeply after 2010, the smoother draws that fall back into the late 2000s and converts a flat-then-break path into a continuous slope. On the unsmoothed series, row (6) of Table 9, the non-London premium is 0.389 in 1997, peaks at 0.415 in 2002, troughs at 0.345 in 2007 and has recovered to 0.369 by 2010—a dip and a recovery rather than a trend, and the same shape as the corresponding line of Figure 4.

Two points of agreement are worth recording. Our estimates and theirs agree that the premium is flat in London throughout, and they agree on the size of the non-London fall over 1997–2019 taken as a whole: rebuilt on our data their series falls from 0.389 to 0.286, and the A-levels line of Figure 4 falls by a similar amount. The disagreement is confined to the timing of the fall within that window. Because their Figure 5 plots the premium against the graduate share rather than against time, and does so on a five-year moving average, it is not built to date the fall, and their text describes the decline as accompanying the 1997–2019 expansion in the graduate share without dating it.

## A.5 The public sector

Figure 21 is the public-sector version of Figure 3. Outside London the public-sector BA/HS differential is not flat before 2010: it rises from $-0.049$ in 1997 to 0 in 2010. It then drifts down, to $-0.023$ in 2019 and $-0.065$ in 2025, against $-0.100$ and $-0.155$ in the private sector. The London series ranges between about $-0.10$ and $+0.10$ with no discernible trend.

Figure 5 in the main text plots the four education $\times$ region wage series for the private sector, where the feature that stands out is the sustained rise in the non-London HS wage after 2014. Figure 22 repeats that exercise on public-sector employees. No such divergence appears: all four series fall together after 2010 and by 2025 lie between roughly 3 and 13 log points below their 2010 level, with the non-London HS series showing none of the sustained upward trend it has in the private sector. One caveat applies to the two London series in particular. The public-sector HS cell in London is small and has been shrinking: it contains 127 wage respondents in 2025, spread over eight five-year age bands, against 2,336 in the non-London regions, and the cell medians that the series is built from therefore rest on as few as seven observations apiece. The year-to-year movement in the London series should be read as sampling variation rather than as a wage trend; the non-London series, which carry the comparison we draw here, are not affected.

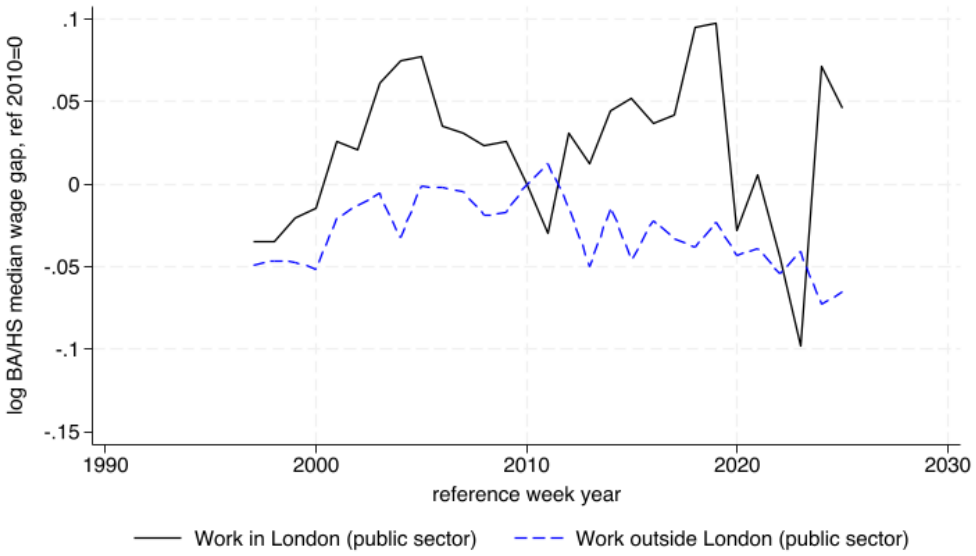


Figure 21: Age-adjusted log BA/HS median wage ratio for public-sector workers, London versus rest of the UK, 1997–2025

*Notes:* LFS public-sector employees (`public`=2) aged 20–59, 1997–2025, by region of workplace; actual-hours wage. The line is the age-adjusted log BA/HS median wage ratio, reference 2010 = 0. The solid line is London and the dashed line the pooled rest of the UK. The London series rests on a small HS cell (see Figure 22) and its year-to-year movement is largely sampling noise.

## A.6 The London post-2010 differential in the public sector

Table 10 repeats the regressions of Table 2 on public-sector employees. The specification is identical bar the sector restriction: the dependent variable is the age-adjusted log median real hourly wage in each education × region × year cell, computed on public-sector employees aged 20–59 over 2001–2020, with region and year fixed effects and standard errors clustered on region; there are 240 cells and 12 clusters in every column, exactly as in Table 2. The table is a robustness exercise and is not part of Table 2, which remains the private-sector table on which the main text is based.

Three differences with the private sector stand out. First, the graduate London × post-2010 coefficient is small and negative rather than zero: −0.018 (standard error 0.004), against 0.005 in the private sector, and −0.025 (0.005) once the NQF shares are added. The NQF shares are only weakly related to public-sector graduate wages: the high-skill share has a coefficient of 0.354 (0.187), significant only at the 10% level and half the private-sector 0.712, and the mid-skill share is insignificant. That is what one would expect if public pay is set on national scales with a fixed London weighting rather than tracking the local graduate labour market, so the link between occupational placement and graduate wages discussed in Section 6 is weaker in the public sector.

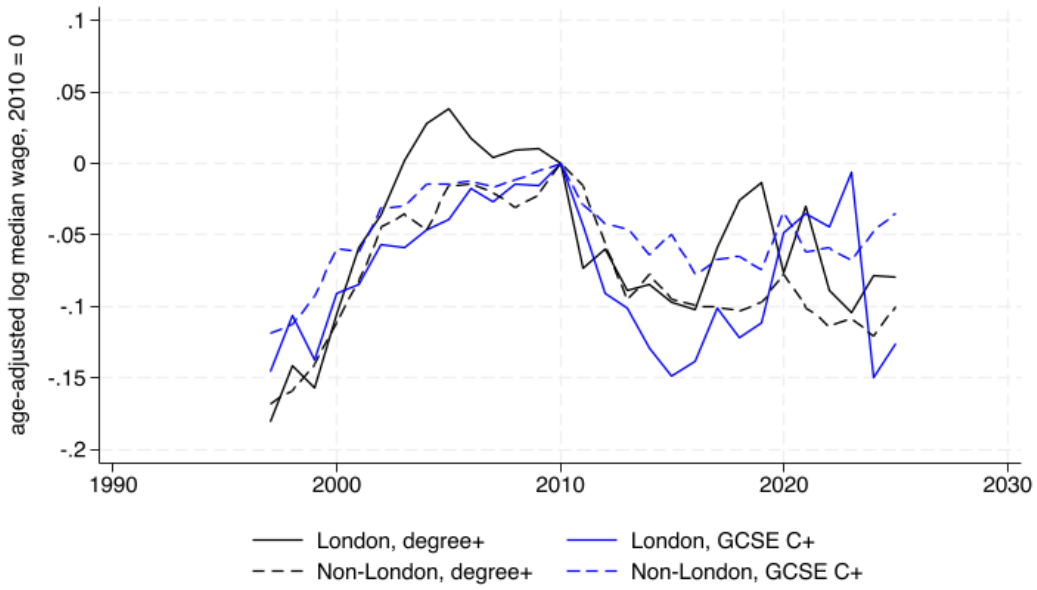


Figure 22: Age-adjusted log median real hourly wage by education group, London versus non-London, public sector, 1997–2025

*Notes:* LFS public-sector employees (`public`=2), aged 20–59, 1997–2025, by region of workplace. The four series are the year fixed effects (reference 2010 = 0) for degree+ and GCSE C+ workers in London and non-London, estimated by OLS on age-band cell medians; the hourly wage is measured using actual hours. The vertical axis runs to −0.20 rather than to −0.15 as in Figure 5, because the public-sector series reach lower in the late 1990s. The London public-sector cells are very thin, and the London GCSE C+ cell is the binding constraint: it contains 127 wage respondents in 2025—7 to 24 per five-year age band—against 2,336 for non-London and 417 for London's private sector, and it has fallen steadily from 719 in 1997 and 438 in 2010. The London degree+ cell is better populated (496 in 2025, 11 to 93 per age band). Much of the year-to-year movement in the solid London GCSE C+ series is therefore sampling noise, and it should not be read as a wage trend.

Second, the HS London penalty is present in the public sector and is statistically significant, but it is roughly half the private-sector size: $-0.036$ (0.008) against $-0.071$. As in the private sector, the estimate for workers below the HS level is more negative than for the HS group; these results are available from the authors. The post-2010 catch-up of non-London wages therefore runs in the same direction in both sectors but is much weaker in the public one. This is consistent with the minimum-wage account of Section 7: Figure 26 shows that the floor never approaches the public-sector HS median, so the part of the private-sector gap that the minimum wage accounts for has no public-sector analogue, while whatever part reflects broader regional labour-market movements would be expected to show up in both. Third, and as in Table 2, adding the NQF shares barely moves the HS coefficient, to $-0.035$; the one NQF share that is significant anywhere in the public-sector table is the high-skill share in the graduate column, and then only at the 10% level.

Two caveats need stating plainly. The first is that the public-sector age cells are smaller than their private-sector counterparts: the graduate and high-school age cells have medians of 38 and 41 wage observations, against 41 and 118.5 in the private sector, although the cells are weighted by size in the age adjustment. The second is that with only 12 clusters, cluster-robust standard errors are biased downward, so the stars in this table and in Table 2 alike should be read as optimistic.

Table 10: Age-adjusted log median wage: London post-2010 differential with NQF occupation controls, public-sector employees

| | Graduates | | High school | |
|---|---|---|---|---|
| | Base | +NQF | Base | +NQF |
| London × post-2010 | -0.018*** | -0.025*** | -0.036*** | -0.035*** |
| | (0.004) | (0.005) | (0.008) | (0.007) |
| NQF high-skill share | | 0.354* | | 0.359 |
| | | (0.187) | | (0.210) |
| NQF mid-skill share | | 0.236 | | 0.119 |
| | | (0.168) | | (0.149) |

Standard errors in parentheses

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

*Notes:* As Table 2 in every respect except the sector: public-sector employees (`public`=2, `econactivity`=1) aged 20–59, 2001–2020. The dependent variable is the age-adjusted log median wage of the education–region–year cell, built from region × five-year age band × year cell medians exactly as in Table 2. All regressions include region and year fixed effects, are unweighted, and cluster standard errors by region (12 clusters); $N = 240$ in every column. Column "Base" includes *London × post-2010* only (London and post-2010 main effects absorbed by fixed effects). Column "+NQF" additionally controls for the share of the education–region–year cell's workers in NQF Level 6+ occupations and the share in NQF Level 3–5 occupations, measured on the same public-sector 20–59 sample as the dependent variable. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

## A.7 Occupation-band wages in the LFS

Figure 8 in the main text tracks real wages by NQF occupation band in the ASHE. Figure 23 is the LFS counterpart, computed on the same private-sector employee sample as our main LFS education plots in Figure 5, with one deliberate difference: it does not select on age. Every other exhibit in this paper imposes an age restriction of some kind – ages 20–59 for the main series, and narrower or older bands elsewhere – whereas this figure imposes none. That is because it is built to be read against the ASHE series, and the published ASHE tables have no age concept at all: they cover employees on adult rates of pay, which is a pay-rate rather than an age criterion. We therefore leave the LFS counterpart unrestricted so that the two are comparable. It tells the same story. The Level 6+ series in London and outside London, and the Level 3–5 series in London, move closely together and none of them has grown much in real terms: on a 2010 base of 100 they stand at 99.7, 99.3 and 103.5 respectively in 2025. The non-London Level 3–5 series is again the outlier, standing 11.0% above its 2010 level in 2025 and 12.2% above it at its 2021 peak. The middle-skilled wage outside London is thus an outlier in the occupation data in the same way as the HS wage outside London is in the education data, which is what we would expect if the two are being driven by the same force.

## A.8 The minimum-wage counterfactual in levels

Figure 12 in the main text plots the non-London minus London HS wage gap, actual and with the real minimum wage frozen at its 2010 level. Figure 24 shows the two level series that underlie it. Freezing the minimum wage shifts both series down, because the year effects absorb whatever the minimum wage did on average and only the cross-region difference is identified; what matters is that the downward shift is larger outside London, where the same national increase is a larger share of the 2010 base median. It is that difference in the shifts which produces the 0.103 against 0.014 log-point comparison quoted in Section 7.

## A.9 The minimum-wage gap bands in real terms

Figure 10 groups regions by the *nominal* gap between the private-sector HS median wage and the NMW. Because the band thresholds are fixed in cash terms, a given gap represents progressively less purchasing power as prices rise, so regions drift back into the wider bands even when the minimum wage has not receded relative to the median. Figure 25 repeats the exercise with both the median and the NMW measured in constant 2025 prices.

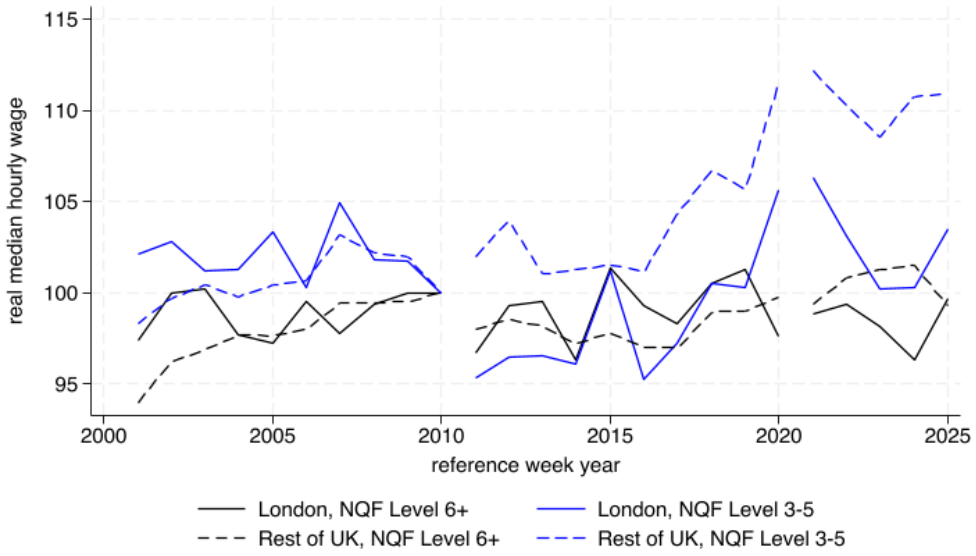


Figure 23: Real median hourly wage by NQF occupation band, London versus non-London, private sector, LFS

*Notes:* LFS, 2001–2025, private-sector employees of all ages reporting a main-job occupation, by region of workplace, income-weighted; the actual-hours wage. Real median hourly wage by the NQF skill level required of the occupation (high-skilled Level 6+ and mid-skilled Level 3–5), in London and non-London, each series indexed to 100 in 2010. Each line is drawn as three separate segments, breaking at the SOC reclassifications (2010/11 and 2020/21), and should be read segment by segment; the 2010 base is common to all three. Occupations are mapped to NQF bands from SOC2000 (2001–2010), SOC2010 (2011–2020) and SOC2020 (2021–2025); the SOC2020 mapping is the same `soc2020_rqf_lookup` crosswalk used for the occupation-share figures.

Running the same nominal exercise on public-sector medians gives a flat picture, shown in Figure 26: the minimum wage never approaches the public-sector HS median. Every one of the twelve regional medians is more than £3 above the NMW in every year to 2016, at most three regions are ever in the £2–3 band, and no region's median comes within £1 of the floor in any year.

# B Comparing the LFS and ASHE Wage Distributions

The concern motivating this appendix is that the post-2010 behaviour of lower-education wages in the LFS might be an artefact of the survey's deteriorating quality rather than a genuine feature of the labour market. Response rates to the LFS fell from just over 50% for new entrants in 2019 to around 30% by 2021 and have stayed near that level through 2025Q4 (Office for National Statistics, 2026c), and a non-trivial share of recorded hourly wages fall below the statutory minimum. We therefore benchmark the LFS wage distribution against the Annual Survey of Hours and Earnings (ASHE). ASHE is an employer-reported,

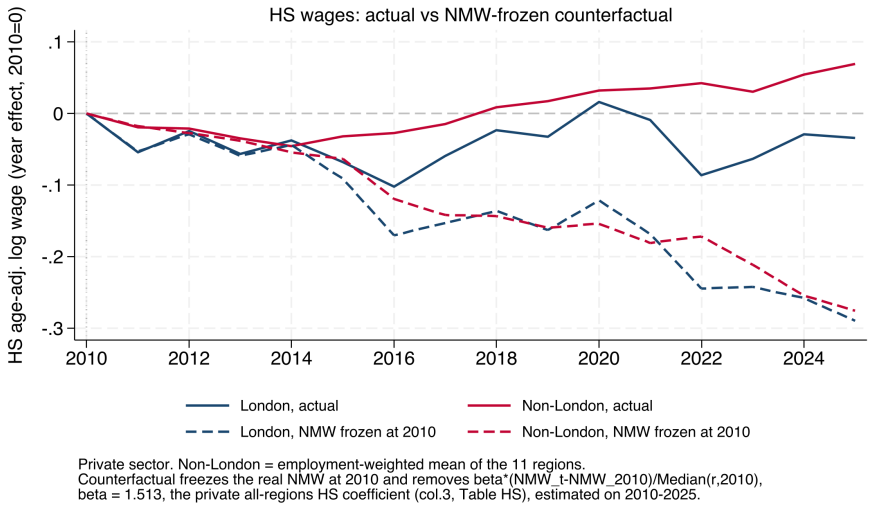


Figure 24: Actual and minimum-wage-frozen counterfactual HS wage, London and non-London, private sector

*Notes:* LFS private-sector HS sample. Age-adjusted log HS wage (year effect, 2010 = 0) for London and non-London: the solid line is the actual series and the dashed line the counterfactual with the real minimum wage frozen at its 2010 level, as in equation (3) with $\beta = 1.513$ (column (3) of Table 3). Non-London is the employment-weighted mean of the eleven regions.

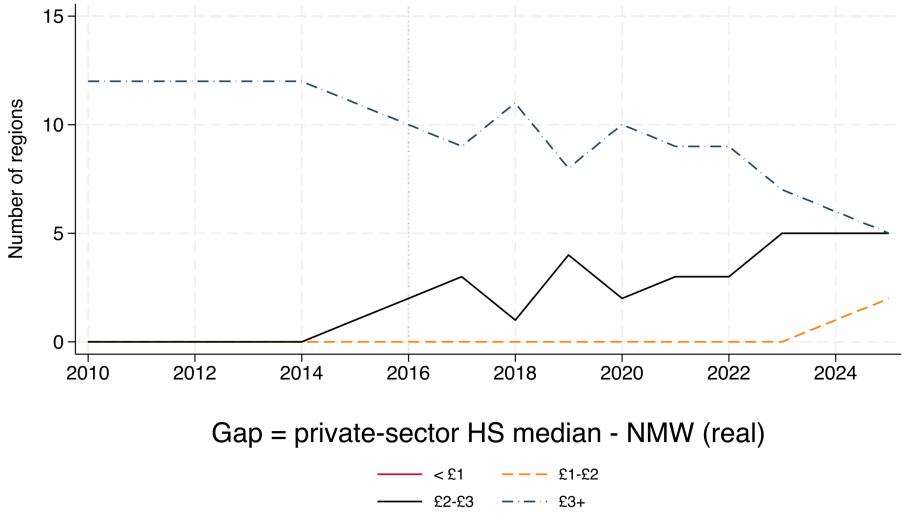


Figure 25: Number of regions by the real gap between the private-sector HS median wage and the minimum wage, by year

*Notes:* As Figure 10 in every respect—LFS HS wage medians by region for 20–59 year olds in the private sector, by region of workplace, and the statutory NMW—except that both the regional median and the NMW are deflated to constant 2025 prices before the gap is taken. The twelve UK regions (including London) are grouped by that gap into four bands (<£1, £1–2, £2–3 and £3+). Because the gap is a difference rather than a ratio, the fixed £ thresholds are not scale-invariant, so the deflation base matters; versions rebased to 2010 and to 2015 prices are produced by the same do-file and tell the same story.

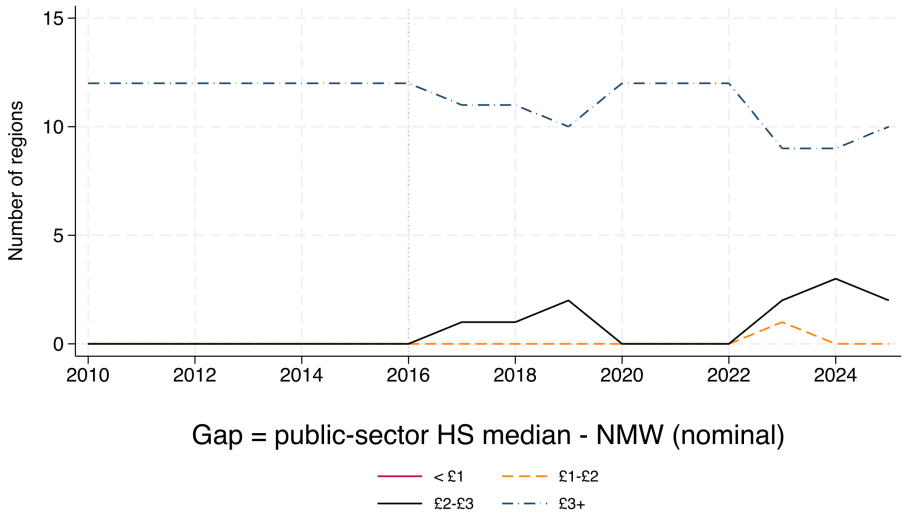


Figure 26: Number of regions by the nominal gap between the public-sector HS median wage and the minimum wage, by year

*Notes:* As Figure 10 in every respect except that the regional HS medians are computed on public-sector employees (`public`=2) rather than private-sector ones: LFS HS wage medians by region for 20–59 year olds, by region of workplace, and the statutory NMW, both in nominal (contemporaneous) terms. The twelve UK regions (including London) are grouped by the gap into four bands (<£1, £1–2, £2–3 and £3+); each line gives the number of regions in each band. HS = GCSE grade C or above without a degree (EDU 2).

administrative survey—an approximately 1% sample of employee jobs drawn from HMRC PAYE records—and so is not subject to the household-survey response problems that afflict the recent LFS. Its published tables cover "employees on adult rates whose earnings for the survey pay period were not affected by absence."[26] Because ASHE contains no measure of education, we cannot compare the graduate premium directly; instead we compare the shape of the wage distribution and, in particular, the behaviour of its lower tail.

We compare ASHE against three LFS hourly-wage measures. The first is `HOURPAY`, the gross hourly-pay variable derived by the ONS: gross weekly pay divided by the sum of basic usual hours and *paid* overtime hours.[27] The second is our own *hourlywage*, gross weekly pay divided by total *actual* hours worked in the main job in the reference week (the LFS variable `TTACHR`); this is the measure used throughout the paper. The two therefore differ in two respects: `HOURPAY` uses *usual* (contractual) hours and counts overtime only when it is paid, whereas our measure uses the hours *actually* worked and counts all of them, paid

[26]ONS, *Guide to interpreting Annual Survey of Hours and Earnings (ASHE) estimates*. "Adult rates" is a pay-rate, not an age, criterion: it excludes employees paid a reduced (junior, development or apprentice) rate, "most of whom will be young people," but—because the minimum wage is only a floor—a young worker paid the ordinary adult rate is included. There is no age threshold.

[27]LFS User Guide, Volume 4 (LFS Standard Derived Variables): `HOURPAY` = `GRSSWK`/(`BUSHR`+`POTHR`), where `GRSSWK` is gross weekly pay in the main job, `BUSHR` is basic *usual* hours, and `POTHR` is *usual* hours of paid overtime (`HOURPAY` = `GRSSWK`/`BUSHR` when the worker reports no paid overtime).

or unpaid. In a normal week with no unpaid overtime the two coincide; they part company when actual hours diverge from usual hours—through leave, sickness or extra hours in the survey week—or when workers do unpaid overtime. Both are ratios of a pay measure to an hours measure and so inherit any error in reported hours. The third is the *reported fixed hourly rate*, available for the roughly one third of employees who report being paid by the hour; it is reported directly and needs no hours denominator, but the hourly-paid are a lower-paid, selected group, so it is informative only near the bottom of the distribution. Since ASHE covers employees on adult rates, the natural LFS comparison sample restricts by age. A 25+ cut is the LFS convention, but young workers paid the adult rate appear in ASHE, so a lower cut of 21+ is arguably a closer match; we show both the 21+ and 25+ samples throughout so that the age cut can be assessed.

## B.1 Snapshot comparisons, 2015, 2019 and 2025

Figure 27 plots the 10th, 25th, 50th, 75th and 90th percentiles of gross hourly pay for the UK as a whole, comparing ASHE with the three LFS measures, for 2015, 2019 and 2025 and for the 21+ and 25+ samples. Three features stand out. First, from the median upward the LFS and ASHE distributions are very close, and the actual-hours *hourlywage* that the paper uses tracks ASHE most closely of the three LFS measures: in 2025 the ASHE median is £17.96, against £18.03 for the 25+ hourlywage and £17.30 for the 21+ hourlywage, and the near-coincidence holds at the 75th and 90th percentiles too (for the 25+ sample the hourlywage 90th percentile, £37.5, is within a pound of ASHE's £37.2). Moving from the 25+ to the 21+ sample lowers the LFS percentiles, because the added 21–24-year-olds are lower paid; the 25+ sample is the closer match to ASHE's level. Second, the one place the LFS falls clearly short of ASHE is the 10th percentile, where the derived LFS measures sit around £11 against ASHE's £12.4 (25+ sample), a gap that has grown since 2015. Third, the reported fixed hourly rate behaves differently at the bottom: its 10th percentile (£12.0 in 2025) is close to ASHE and well above the derived measures, even though it falls below them higher up because the hourly-paid are a lower-paid group. That the directly-reported rate agrees with ASHE at the bottom while the derived wages do not points to mismeasured hours as the source of the low-tail gap.

The lower-tail discrepancy is sharpest in low-paying occupations, where pay sits close to the minimum. Figure 28 repeats the exercise for elementary occupations, aged 25+, in 2019 and 2025, and here the derived LFS measures are implausibly low at the bottom. In 2025 the ASHE 10th percentile is £11.5, whereas the derived LFS hourlywage puts it at about £9.1—below the £12.21 adult minimum wage in force from April 2025. The reported fixed hourly rate, by contrast, gives an elementary 10th percentile of about £11.5, close to ASHE. That

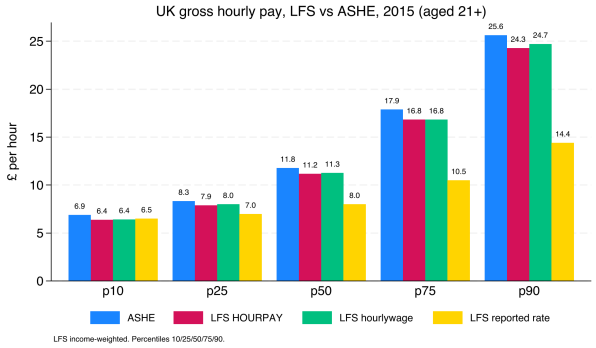


(a) 2015, aged 21+

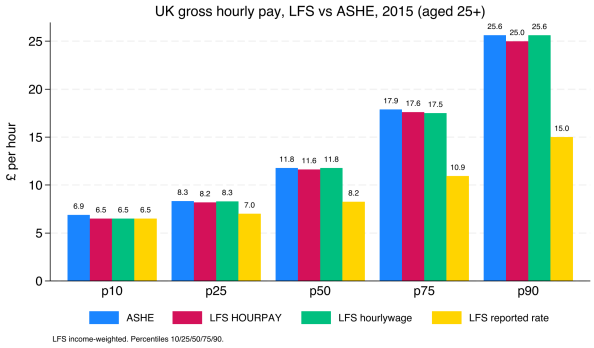


(b) 2015, aged 25+

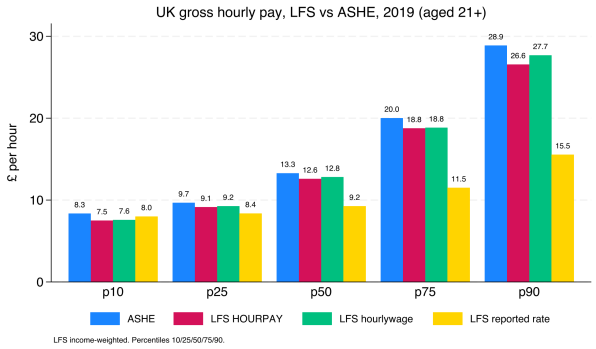


(c) 2019, aged 21+

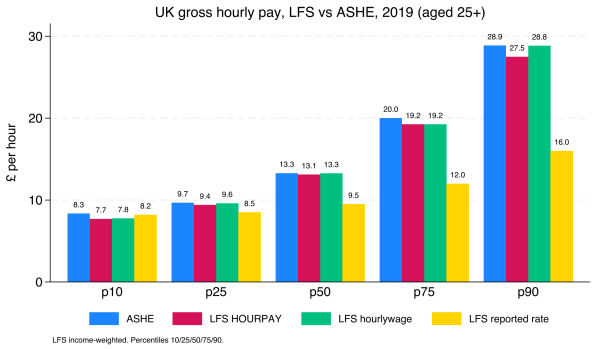


(d) 2019, aged 25+

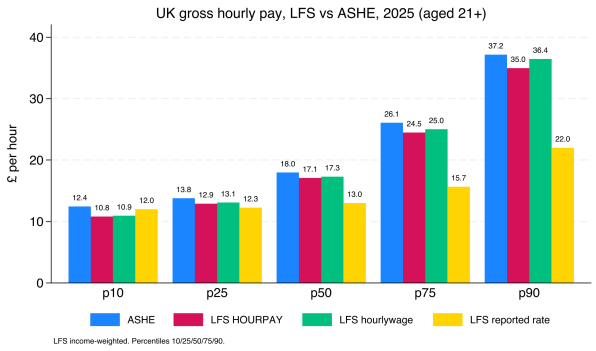


(e) 2025, aged 21+

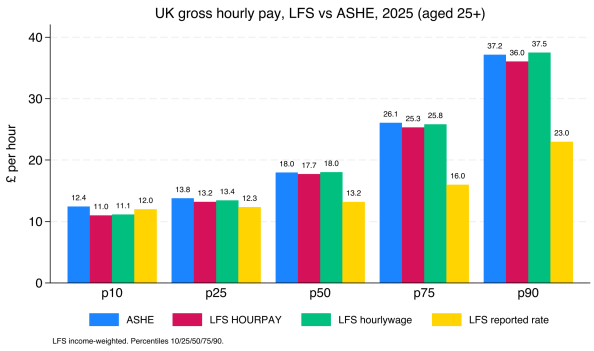


(f) 2025, aged 25+

Figure 27: UK gross hourly pay distribution, LFS versus ASHE

*Notes:* Snapshots for 2015, 2019 and 2025. Each group of bars is a percentile (the 10th, 25th, 50th, 75th and 90th); within each group the bars are ASHE, the LFS `HOURPAY` measure, the LFS hourlywage (actual hours) and the LFS reported fixed hourly rate. The left column is aged 21+ and the right column aged 25+. LFS estimates are income-weighted, and ASHE 2025 is provisional.

the directly-reported rate agrees with ASHE while the *derived* wage falls below the legal floor—which the employers ASHE records are bound by—shows the gap to be measurement rather than genuinely sub-minimum pay: it is the signature of mismeasured hours in the derived-wage denominator, concentrated where wages are low and hours short or irregular. This is on the 25+ sample that best matches ASHE's adult-rate coverage, so composition is not the explanation. The same ranking is visible, more mildly, in 2019, when the minimum wage stood lower relative to the median.

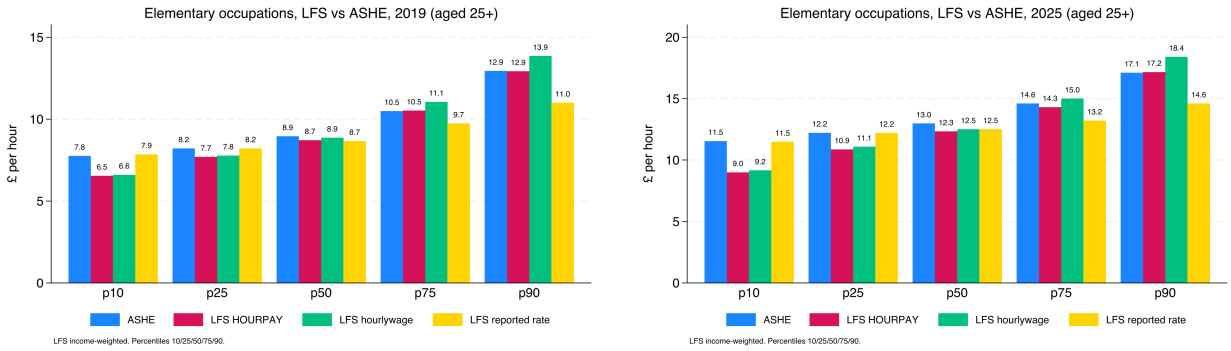


(a) 2019, aged 25+

(b) 2025, aged 25+

Figure 28: Elementary-occupation gross hourly pay distribution, LFS versus ASHE

*Notes:* Snapshots for 2019 and 2025, aged 25+. Within each percentile group the bars are ASHE, the LFS `HOURPAY` measure, the LFS hourlywage (actual hours) and the LFS reported fixed hourly rate, as in Figure 27.

## B.2 Trends, 2001–2025

Figure 29 plots the median gross hourly wage from the two sources for every year from 2001 to 2025, in real 2025 prices, for the 21+ and 25+ LFS samples.[28] The correspondence is remarkably good. For the 25+ sample the LFS and ASHE medians are almost indistinguishable across the whole period—the ratio of the two stays within about one per cent of unity in most years—the main exception being a transitory divergence in 2020–22, when pandemic disruptions to working time (furlough) temporarily depressed *actual* hours and so inflated the actual-hours LFS wage; the two series reconverge by 2025 (£18.2 against £18.1). For the 21+ sample the LFS median runs a steady three to five per cent below ASHE (£17.5 against £18.1 in 2025), because it adds lower-paid 21–24-year-olds, but it moves in the same close parallel. Either way, the deterioration in LFS response rates has not driven the LFS

[28]The ASHE median (all employees, UK) is the published series; the LFS median uses the actual-hours wage, income-weighted. Both are deflated to 2025 prices with the GDP deflator, so the comparison is unaffected by the choice of base year.

and ASHE medians apart, which is what matters for the median-based education-premium series in the body of the paper.

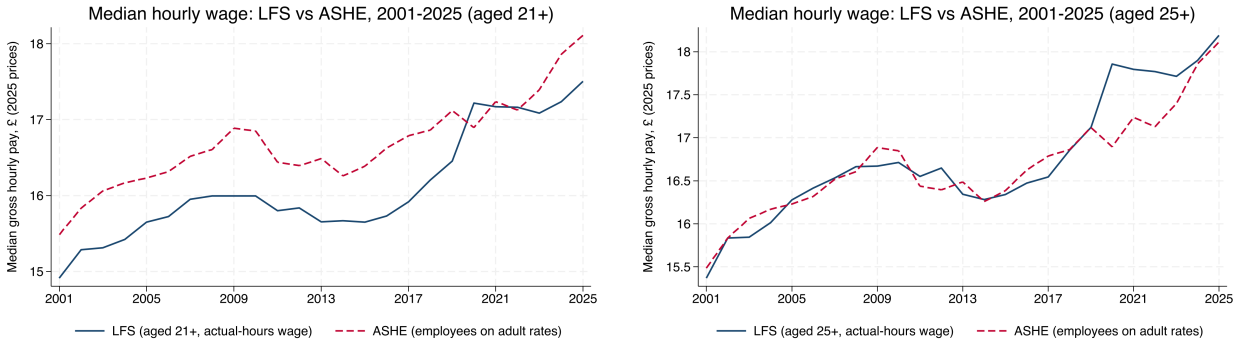


(a) aged 21+ (b) aged 25+

Figure 29: Median gross hourly pay, LFS versus ASHE

*Notes:* Annual, 2001–2025, in real 2025 prices. The LFS series is the actual-hours wage, income-weighted; the ASHE series covers all employees on adult rates (2025 provisional).

The median comparison holds region by region, not just for the UK as a whole. Figure 30 plots the LFS (25+) and ASHE median separately for each of the twelve regions, back to 2001. The two series move together in every region, though not exactly: in most regions they stay within a few per cent of each other throughout, while in London the LFS median runs as much as £1.40—nearly six per cent—below ASHE through the late 2000s and early 2010s before the two converge after 2015. Across the twelve regions in 2025 the LFS median differs from ASHE by 1.9 per cent on average and by at most 3.5 per cent, and no region drifts away persistently, so the close national correspondence is not masking regional gaps large enough to unsettle the paper's median-based comparisons. The figure also settles a specific concern—that a low-wage region's median might sit down among the mismeasured low wages at the bottom of the national distribution. It does not. Even the lowest-paid region's median lies far above the UK 10th percentile: in 2025 the lowest regional median (East Midlands, about £16) exceeds the UK 10th percentile by nearly £5 and sits some £2.6 above even the UK 25th percentile, and the same ordering holds in every year. The regional medians on which the paper's London-versus-rest comparisons rest therefore lie well clear of the part of the distribution where the LFS and ASHE part company.

Figure 31 plots the 10th and 25th percentiles from both sources, together with the real value of the adult minimum wage. The 25th percentiles move in close parallel and all but coincide—the LFS runs about two to three per cent below ASHE by 2025—so the low-tail discrepancy is essentially confined to the tenth percentile. There the two diverge. Until the early 2010s the LFS 10th percentile sat above the then-low real minimum wage; as the minimum wage rose steeply after 2015 it drew level with and then overtook the LFS 10th

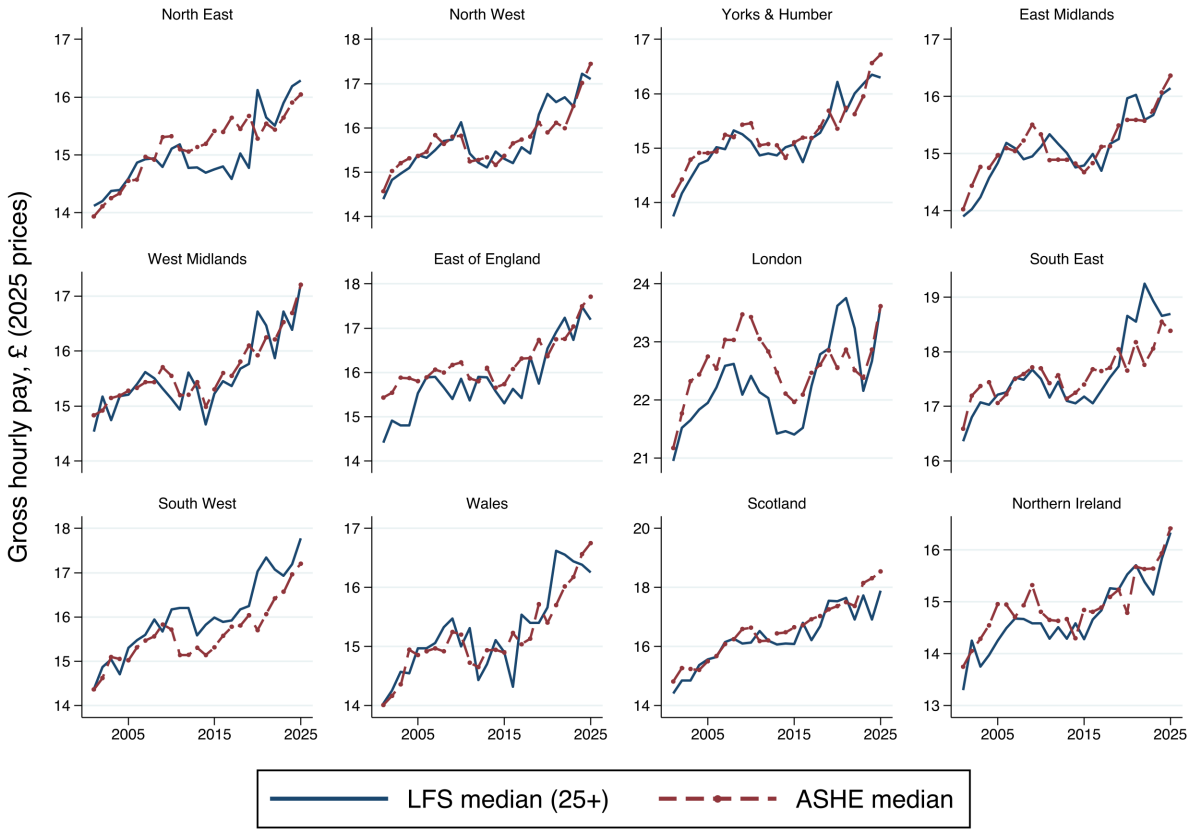


Figure 30: Median gross hourly pay by region, LFS (aged 25+) versus ASHE

*Notes:* Real 2025 prices, 2001–2025. Each panel is one of the twelve UK regions; the solid line is the LFS median (actual-hours wage, income-weighted, by region of workplace) and the dashed line the ASHE regional median (employees on adult rates).

percentile, which by 2025 sits below the real adult minimum (£11.0 for 21+, £11.3 for 25+, against £12.0). The ASHE 10th percentile, available from 2015, instead stays well above the minimum (£12.5 in 2025), and its gap over the LFS 10th percentile widens over 2019–2025. Using the 25+ rather than the 21+ sample lifts the LFS 10th percentile only slightly. That the divergence appears only at the very bottom, only in the LFS, and only once the minimum wage rises—while the employers whom ASHE records are legally bound by that minimum—again points to measurement error in the derived low wages rather than to a genuine collapse of pay beneath the floor.

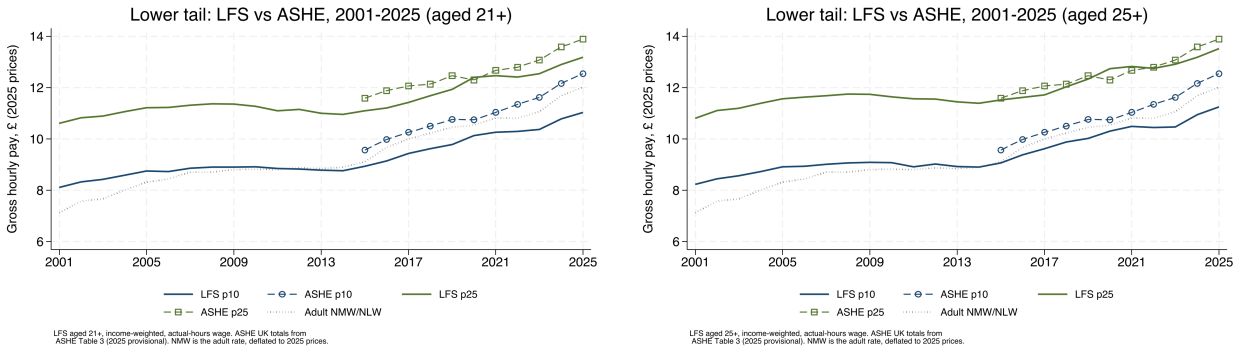


(a) aged 21+ (b) aged 25+

Figure 31: Lower-tail gross hourly pay, LFS versus ASHE

*Notes:* The 10th and 25th percentiles, annual 2001–2025, in real 2025 prices, with the real adult minimum wage overlaid. The LFS series is the actual-hours wage, income-weighted. ASHE UK totals are available for 2015–2025 (2025 provisional).

## B.3 Implications

Taken together, these comparisons support treating the LFS wage patterns in the body of the paper as genuine. Measured on the paper's actual-hours basis, the LFS median matches ASHE almost exactly for the 25+ sample, and runs a few per cent below it in close parallel for the 21+ sample, throughout 2001–2025; and the two sources agree across the upper half of the distribution. The median-based, trend-focused education-premium series are therefore not distorted by the LFS's response-rate problems. The one place the two sources genuinely disagree is the extreme lower tail in the years since the minimum wage began rising steeply—where derived hourly wages are most exposed to hours mismeasurement and where the LFS retains low-paid workers that ASHE's adult-rate coverage excludes. Because the high-school median wage on which our main results rest lies far above this problematic region, the post-2010 rise in the non-London high-school wage cannot be dismissed as a measurement artefact.

# C Robustness of the Minimum-Wage Spillover Regressions

This appendix reports the robustness checks summarised in the main text for the regional minimum-wage spillover regressions—the high-school (HS) wage regression on the minimum-to-median bite (Table 3), with the graduate wage as a placebo (Table 4). We re-estimate the specification over alternative sample windows, and re-run it across the HS wage distribution. Throughout, the specification, weighting and clustering match the headline regression.

**Sample period.** The headline regressions use the 2010–2025 window, with the bite's denominator frozen at the region's 2010 HS median. Table 11 re-estimates the HS regression (private sector, all regions) in four ways: on the two disjoint sub-periods 2010–2015 and 2016–2025, which split the sample at the introduction of the National Living Wage, and on the two pooled single-slope windows 2010–2025 and 2012–2025, which change the share of the sample drawn from the post-2015 period of rapid minimum-wage growth. The four point estimates are close to one another and to the headline estimate: 1.71 on 2010–2015, 1.64 on 2016–2025, 1.51 on 2010–2025 and 1.64 on 2012–2025. Precision, however, differs sharply. The three windows that contain the post-2015 minimum-wage increases are all significant at the 1% level, whereas the 2010–2015 sub-period carries a standard error of 1.85 and is not distinguishable from zero ($p = 0.37$)—which is unsurprising, since the real minimum wage rose by only 3.3% over those years, so there is very little movement in the bite for the regression to use. We would therefore not treat the 2010–2015 column as independent evidence either way; what matters is that the slope does not change materially once the identifying variation is present. The graduate placebo (Table 12) is small and statistically insignificant in every window.

Table 11: HS wage on the bite, by sample window

| | Disjoint sub-periods | | Pooled windows | |
|---|---|---|---|---|
| | 2010–2015 | 2016–2025 | 2010–2025 | 2012–2025 |
| $\text{NMW}_t/\text{Median}_{r,2010}$ | 1.713 | 1.642*** | 1.513*** | 1.641*** |
| | (1.849) | (0.400) | (0.144) | (0.199) |
| Observations | 72 | 120 | 192 | 168 |

Standard errors in parentheses

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

*Notes:* Private sector, all regions; outcome is the age-adjusted median HS year effect; bite = $\text{NMW}_t/\text{Median}_{r,2010}$; region and year fixed effects; weighted by cell employment; SE clustered by region. The first two columns fit the slope on *disjoint* sub-periods (2010–2015 and 2016–2025); the last two are pooled single-slope windows (2010–2025 and 2012–2025). The 2010–2015 column has little identifying variation: the real minimum wage rose only 3.3% over those years.

Table 12: Graduate (BA) wage on the bite, by sample window (placebo)

| | 2010–2025 | 2012–2025 | 2016–2025 |
|---|---|---|---|
| $\text{NMW}_t/\text{Median}_{r,2010}$ | 0.142 | 0.138 | 0.165 |
| | (0.469) | (0.400) | (0.311) |
| Observations | 192 | 168 | 120 |

Standard errors in parentheses

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

*Notes:* As Table 11 but for degree-holders, run on the graduate base bite. Region and year fixed effects; weighted by cell employment; SE clustered by region.

**Distribution: which HS wages move?** If the mechanism is a wage floor, the minimum wage should raise the bottom of the HS distribution more than the top. Table 13 re-runs the HS regression at the 25th, 50th and 75th percentiles (excluding London). The estimates are large at the bottom and middle and fall away at the top. In the private sector, our preferred sample, the estimates are 1.31 at the 25th percentile, 1.60 at the median and 0.25 (s.e. 0.36) at the 75th—exactly the shape a wage floor predicts, with compression from below that dies out by the top, where the estimate is about a sixth of the median and not distinguishable from zero. For all workers the corresponding figures are 0.91, 0.62 and −1.13. We would not read that negative all-worker estimate at the 75th percentile as an effect of the floor; it belongs with the other all-worker anomalies we attribute to public-sector composition—compare the negative all-worker graduate coefficients in Table 4—and it is absent from the private-sector

sample. Note that on the all-worker sample the median estimate is itself imprecise (0.62, s.e. 0.38), so the shape across percentiles is clearly identified only in the private sector. Were we instead capturing a broad force such as productivity growth, the 75th percentile would be expected to move as well. The occupation-level data (Section 7.2.2) deliver the same shape far more sharply.

Table 13: Effect across the HS wage distribution (p25, p50, p75)

| | All workers (excl. London) | | | Private sector (excl. London) | | |
|---|---|---|---|---|---|---|
| | p25 | p50 | p75 | p25 | p50 | p75 |
| $\text{NMW}_t/\text{Median}_{r,2010}$ | 0.911* | 0.620 | -1.134* | 1.306*** | 1.600*** | 0.250 |
| | (0.448) | (0.376) | (0.541) | (0.200) | (0.227) | (0.355) |
| Observations | 176 | 176 | 176 | 176 | 176 | 176 |

Standard errors in parentheses

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

*Notes:* Outcome is the age-adjusted HS year effect at the indicated percentile, region×year, 2010–2025, excluding London. Regressor is the HS base bite. Region and year fixed effects; weighted by cell employment; SE clustered by region.

### Relative supply with no other controls

The decomposition in Table 7 attributes essentially nothing to the relative supply of graduates. Because that decomposition also conditions on the industry and occupation mix, a reader might reasonably worry that relative supply operates *through* composition and is absorbed by those blocks. Table 14 removes every other control, so that relative supply is not conditioned on any variable through which it might plausibly act.

## D Normalising the Minimum-Wage Exhibits to 2014

Every exhibit in the body of the paper measures the non-London−London high-school wage gap from **2010**, the first year of our estimation window. An alternative is to measure it from **2014**, on the grounds that 2014 is the year in which HS wages inside and outside London actually part ways: on our private-sector, age-20–59 sample the gap is at its *lowest* in 2014 (−0.008 log points, against 0 by construction in 2010 and +0.103 in 2025). Measuring from the trough rather than from an arbitrary window start makes the post-2014 divergence the object of interest, and it turns out to sharpen the paper's main conclusion considerably. This appendix reports the four affected exhibits on that basis so the choice can be made deliberately. We have not adopted it in the body, for a reason we set out at the end.

Table 14: The relative supply of graduates with no other controls, private sector.

| | Supply only | | Supply + bite | |
|---|---|---|---|---|
| | (1) | | (2) | |
| | Log pts | % of gap | Log pts | % of gap |
| Baseline gap $\theta_{\text{base}}$ | 0.103 | 100.0 | 0.103 | 100.0 |
| BA/HS supply | -0.000 | -0.4 | -0.001 | -0.5 |
| | (0.000) | | (0.000) | |
| NMW bite | — | — | 0.097 | 93.9 |
| | | | (0.012) | |
| Total explained | -0.000 | -0.4 | 0.097 | 93.4 |
| | (0.000) | | (0.012) | |
| Unexplained $\theta_{\text{full}}$ | 0.104 | 100.4 | 0.007 | 6.6 |
| | (0.016) | | (0.016) | |

*Notes:* As Table 7, but the only covariates are the BA/HS worker ratio in column (1) and that ratio together with the minimum-wage bite in column (2). The industry, occupation, gender/public and immigration blocks are omitted entirely, so relative supply is not conditioned on any variable through which it might plausibly operate. Standard errors in parentheses: wild-cluster bootstrap (Webb weights, 9999 replications) for the contributions and the total; clustered by the twelve regions for $\theta_{\text{full}}$. Reproduced by `syntax/composition adjusted wage over time.do` (Section 2 Step 9, block (C2)).

## D.1 The minimum wage on its own

Figures 32 and 33 are the 2014-normalised counterparts of Figures 24 and 12. The construction is unchanged—the counterfactual still removes $\beta(\text{NMW}_t - \text{NMW}_{2010})/\text{Median}^{HS}_{r,2010}$ with $\beta = 1.513$—and only the location of the zero moves: every series is shifted to pass through zero in 2014.

One consequence deserves emphasis, because it means the change is not purely cosmetic. Re-referencing both the actual and the counterfactual series to 2014 subtracts the same 2014 value from each, so for any year $t$

$$(y_t - y_{2014}) - (\tilde{y}_t - \tilde{y}_{2014}) \;=\; \beta\big(\text{NMW}_t - \text{NMW}_{2014}\big)\Delta(1/\text{Median})\,,$$

that is, from 2014 onwards the vertical distance between the two lines is exactly what a counterfactual freezing the real minimum wage at its *2014* level would deliver. Normalising to 2014 therefore silently re-bases the counterfactual as well as the axis. Only the pre-2014 segment still reflects the 2010 freeze, and there the real minimum wage barely moved, so little is lost—but the two figures are answering slightly different questions.

On this basis the actual gap reaches 0.112 log points by 2025 while the minimum-wage-frozen counterfactual reaches 0.024, so the rising floor accounts for roughly four-fifths of the

post-2014 divergence.

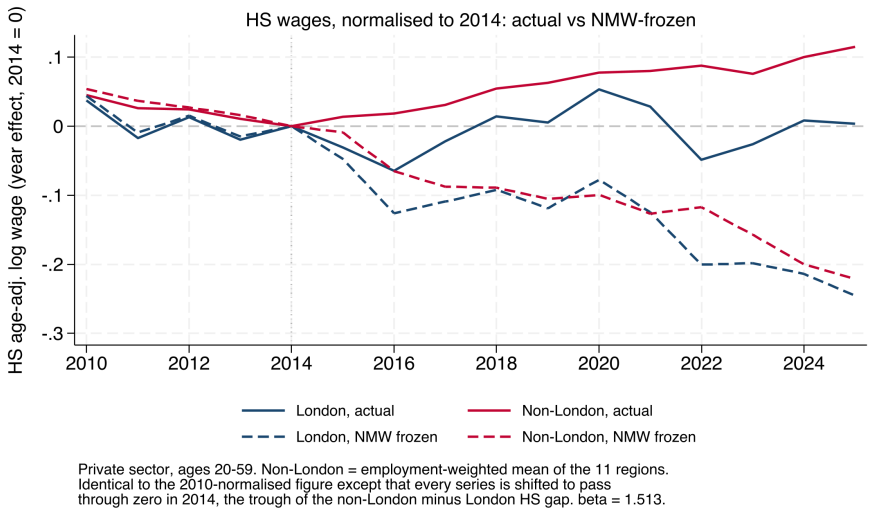


Figure 32: Actual and minimum-wage-frozen counterfactual HS wage, normalised to 2014

*Notes:* As Figure 24 in every respect—LFS private-sector HS sample, employees aged 20–59, $\beta = 1.513$ from column (3) of Table 3, non-London the employment-weighted mean of the eleven regions—except that each of the four series is shifted to pass through zero in 2014 rather than 2010.

## D.2 The decomposition

Table 15 and Figure 34 repeat Table 7 and Figure 16 with the omitted non-London×year interaction moved from 2010 to 2014, so that $\theta_{2025}$ is the 2025 gap measured from the trough. Sample, window, weighting, clustering, covariate blocks, the Gelbach (2016) decomposition and the wild-cluster bootstrap are all identical; only the reference year differs. The bite's own slope is unchanged at $\hat{\beta} = 2.07$, as it must be—dropping a different interaction year does not alter the span of the regressors.

This is where the case for the change is strongest, and it is the point the exercise was proposed to test. Measured from 2010, the composition blocks together take 0.018 of the 0.103 gap, about a sixth of it. Measured from 2014 they take −0.001 of the 0.112 gap: essentially nothing, and if anything the wrong sign. The gender margin falls from 0.009 to 0.006, the public-sector margin from 0.008 to 0.001, and occupation from 0.012 to 0.005, while industry moves further negative. The minimum-wage bite, by contrast, is allocated 0.120 log points, or 107% of a gap that is itself larger. So on the 2014 basis the composition narrative is not merely outweighed—it explains none of the divergence at all.

Table 15: Gelbach decomposition of the 2025 non-London−London HS wage gap measured from 2014, private sector.

| | Composition only | | With NMW bite | |
|---|---|---|---|---|
| | (1) | | (2) | |
| | Log pts | % of gap | Log pts | % of gap |
| Baseline gap $\theta_{\text{base}}$ | 0.112 | 100.0 | 0.112 | 100.0 |
| Industry | -0.005 | -4.2 | -0.014 | -12.2 |
| | (0.010) | | (0.013) | |
| Gender (female share) | 0.006 | 5.2 | 0.006 | 5.4 |
| | (0.002) | | (0.002) | |
| Sector (public share) | 0.001 | 1.1 | 0.001 | 1.0 |
| | (0.001) | | (0.001) | |
| Gender × sector | 0.003 | 3.1 | 0.003 | 2.7 |
| | (0.001) | | (0.001) | |
| Occupation | 0.000 | 0.2 | 0.005 | 4.7 |
| | (0.008) | | (0.006) | |
| Immigration | 0.001 | 0.9 | -0.001 | -1.1 |
| | (0.001) | | (0.000) | |
| BA/HS supply | -0.002 | -1.8 | -0.002 | -1.6 |
| | (0.001) | | (0.001) | |
| NMW bite | — | — | 0.120 | 107.4 |
| | | | (0.018) | |
| Total explained $\theta_{\text{base}} - \theta_{\text{full}}$ | 0.005 | 4.6 | 0.119 | 106.4 |
| | (0.012) | | (0.015) | |
| Unexplained $\theta_{\text{full}}$ | 0.107 | 95.4 | -0.007 | -6.4 |
| | (0.020) | | (0.019) | |
| *Memo: NMW-bite slope in column (2)* | | | | |
| NMW bite coefficient $\hat{\beta}$ | — | | 2.074* | |
| | | | (0.403) | |
| wild-bootstrap $p$ | — | | 0.056 | |
| wild-bootstrap 95% CI | — | | [-0.078, 4.105] | |

*Notes:* Identical to Table 7 except that the omitted non-London×year interaction is 2014 rather than 2010, so $\theta_{\text{base}} = 0.112$ is the 2025 gap measured from the 2014 trough rather than the 0.103 measured from 2010. Employment-weighted, 2010–2025, employees aged 20–59; the dependent variable and the bite are built on the private-sector HS sample while the composition controls are region×year shares over all HS employees. Standard errors in parentheses: wild-cluster bootstrap (Webb, 9999) for the blocks and the total, clustered by the twelve regions for $\theta_{\text{full}}$ and the memo slope. Reproduced by `composition adjusted wage over time.do`, the `basey(2014)` call.

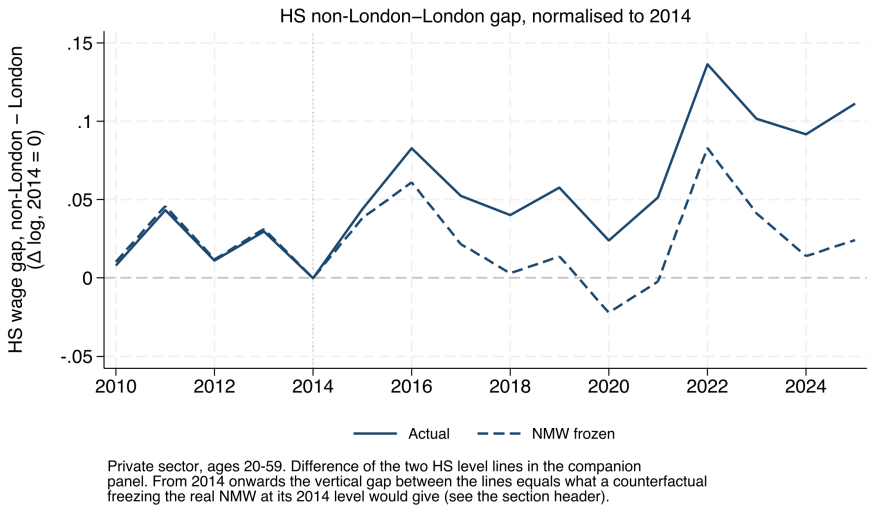


Figure 33: Non-London minus London HS wage gap, actual and counterfactual, normalised to 2014

*Notes:* The difference between the two pairs of lines in Figure 32, so that positive values mean non-London above London and 2014 = 0. From 2014 onwards the vertical gap between the solid and dashed lines equals what freezing the real minimum wage at its 2014 level would give, as shown in the text above.

## D.3 Why we have not adopted it

The 2014 normalisation makes the paper's argument stronger, and everything above is reproducible on exactly the same sample and specification as the body. Our hesitation is presentational rather than statistical. The whole of Sections 8 and the minimum-wage section is framed around changes since 2010: the bite is $\text{NMW}_t/\text{Median}^{HS}_{r,2010}$, the frozen median is a 2010 median, the estimation window opens in 2010, and every other figure in the paper is indexed to 2010. Switching the base in the decomposition exercise (and the earlier parts of the paper) to 2014 invites the criticism that we are choosing our sample presentation to emphasise minimum wage effects, since it is the only variable that changes after but not before 2014. Thus, the 2010 basis is the more conservative choice and, as Table 7 shows, it already supports the conclusion; the 2014 basis can, in part, be read as evidence that the conclusion is not fragile to where the clock starts.

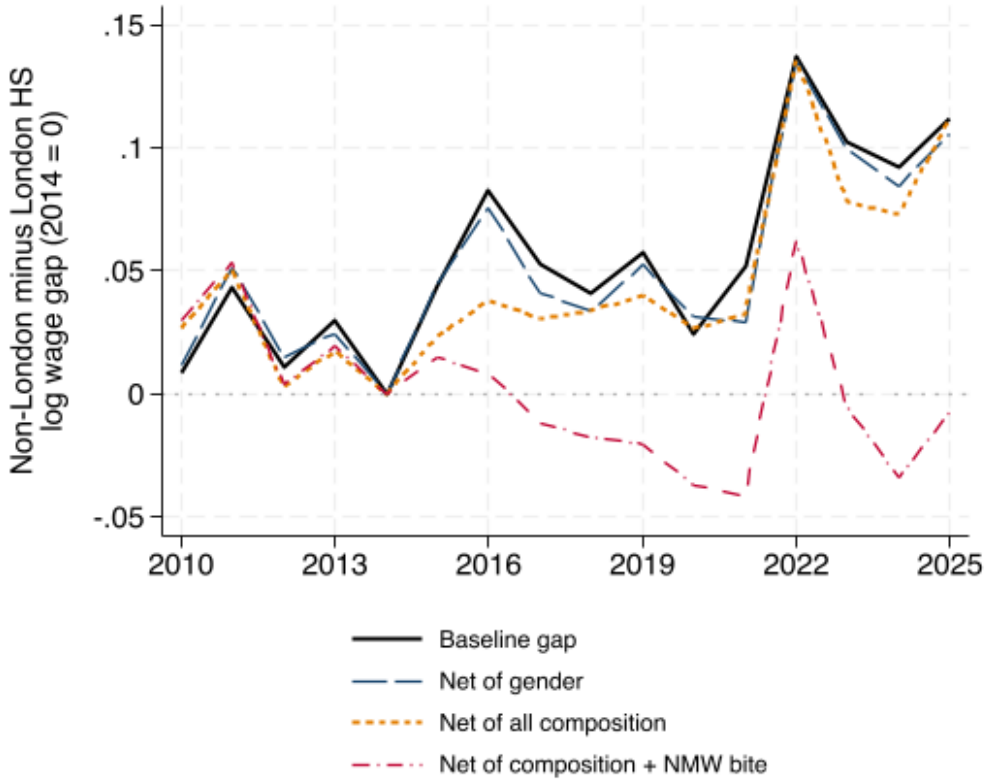


Figure 34: Cumulative decomposition of the non-London HS wage gap, normalised to 2014

*Notes:* As Figure 16 but with every path shifted to pass through zero in 2014. In 2025 the baseline gap is 0.112; removing the gender margin leaves 0.106; removing all of the composition controls leaves 0.113—slightly *above* the baseline, because composition is allocated a small negative contribution on this basis; and removing the minimum-wage bite as well leaves −0.007. Reproduced by `composition adjusted wage over time.do`, the `basey(2014)` call.