



Inequality

The IFS Deaton Review

Spatial disparities across labour markets

**Henry G. Overman
Xiaowei Xu**

An IFS initiative funded by the Nuffield Foundation

Spatial disparities across labour markets

Henry G. Overman (LSE; CEP) and Xiaowei Xu (IFS)¹

Executive summary

- **Spatial disparities in wages and employment rates are large and persistent, although they are smaller today than 20 years ago.** In 2019, average wages in London were £20 per hour, 60% higher than the £13 paid in Scarborough and Grimsby. Employment rates ranged from 66% in Skegness and Louth to 90% in Harrogate.
- **Much of the difference in wages between areas is driven by wages for the higher-paid.** In 2019, the wages of the top 10% of earners in London were £37, over 80% higher than the £20 paid to the top 10% of earners in Scarborough. The wages of the bottom 10% of earners were similar everywhere at around £8–£9.
- **Spatial disparities in labour market outcomes largely reflect the concentration of high-skilled workers, who would have better labour market outcomes wherever they live.** The share of adults with degrees ranges from 15% in Doncaster to 55% in Brighton. High-skilled workers tend to work in better-performing labour markets, which further magnifies individual labour market advantages. At least 64% and up to 90% of differences in average wages across areas can be attributed to differences in the types of people who work in different places.
- **The spatial concentration of high-skilled workers is driven by differences in the demand for, and supply of, skills and the interaction between the two.** The demand for high-skilled workers is spatially concentrated and possibly becoming more so. There are large differences in educational attainment across areas, that are exacerbated by patterns of graduate migration.
- **The self-reinforcing interaction between demand and supply is particularly pronounced in the highest-wage areas and at the upper end of the wage distribution.** The result is that the highest-paid jobs are concentrated in London and a handful of other areas.
- **Places that offer higher earnings also have higher rents, which may entirely offset gains in earnings.** Consistent with this, spatial disparities in well-being are smaller than disparities in labour market outcomes. People in higher-paid places are no happier than those in lower-paid places.
- **Differences in labour market opportunities, costs of living and amenities help explain some of the large differences in age structure across the country.** Young people are increasingly concentrated in London and other cities, which offer good labour market opportunities. Older adults are concentrated in coastal and rural areas, which offer natural amenities and low costs of living.

¹ The authors are grateful to James Banks, Richard Blundell, Angus Deaton and Robert Joyce for their insightful comments over many discussions. This work was produced using statistical data from the Office for National Statistics (ONS). The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. This work uses research data sets which may not exactly reproduce National Statistics aggregates. All errors are the authors' own.

- **Places matter to people but care is needed when thinking about the impact of place-based policies.** The link between individual and spatial disparities is complicated by the fact that some people move, which means that place-based policies do not necessarily end up benefiting the people that they aim to help. In many cases, using a place-based approach to targeting policy will also fail to address within-area disparities or reduce individual inequality.

1. Introduction

Spatial inequalities in the UK are profound and persistent. In 2019, average wages in London were 60% higher than those in Scarborough and Grimsby. Employment rates ranged from 66% in Skegness and Louth to 90% in Harrogate. Around half of working-age adults in London and Brighton had degrees, compared with less than a fifth in places such as Doncaster, Mansfield and Grimsby.

These spatial inequalities have concerned successive governments and currently sit high up the political agenda. The Conservatives under Boris Johnson have made 'levelling up' the core of their programme for government, with Michael Gove, the head of the Department for Levelling Up, Housing and Communities, calling it the 'defining mission of this government'.² Polling shows that inequalities between more and less deprived areas are considered the most serious form of inequality in Britain, and an issue on which there is significant agreement across the political spectrum (Benson et al., 2021).

There are many dimensions of inequality between places – including in living standards, health and educational attainment – and the extent of these disparities varies at different spatial scales. This chapter focuses on spatial disparities in employment and wages across local labour markets. To understand the causes and consequences of these disparities, we explore spatial disparities in education and demographics as well as patterns of graduate migration. We also consider disparities in costs of living and amenities – which people trade off against labour market opportunities when deciding where to live – as well as differences in self-reported well-being.

Several considerations justify a focus on spatial disparities in labour market outcomes. For many households, earnings are the most important source of income and thus a key determinant of living standards. Disparities in employment and wages also matter for disparities in health (Case and Kraftman, forthcoming) and many other socio-economic outcomes. Moreover, people care about spatial disparities in labour market outcomes. In a recent survey of what people want from the levelling-up agenda, improved labour market prospects ('better job opportunities in your area') emerged as the top priority (Carter, 2021).

A key question for any analysis of spatial disparities is whether the area affects individual outcomes. At one extreme, spatial disparities may represent the spatial manifestation of individual inequality, reflecting the fact that individuals with different outcomes live in different places. At the other extreme, local conditions might be the only determinant of differences in individual outcomes and inequality. Understanding this link between individual and spatial disparities is complicated by the fact that, while many people stay close to where they grow up, many others move around.

² <https://www.gov.uk/government/news/ambitious-plans-to-drive-levelling-up-agenda>.

Our analysis shows that the spatial concentration of high-skilled workers and the firms that employ them, and the self-reinforcing nature of spatial differences in the supply of and demand for skills, play a key role in explaining the extent of disparities and their persistence. The concentration of high-skilled workers matters because they have better labour market outcomes wherever they live. Thus, areas where they are concentrated will have better outcomes on average. Where these high-skilled workers are concentrated also matters because some areas generate higher productivity and better labour market outcomes. We show that skilled workers generally work in better-performing labour markets and vice versa. This self-reinforcement is particularly pronounced in the highest-wage areas and at the upper end of the wage distribution. The result is that the highest-paid jobs are concentrated in London and a handful of other areas.

Of course, labour market outcomes are not all that matter. When deciding where to live, people trade off the benefits and costs of different areas – the earnings they expect to make, the cost of living, access to the natural environment, safety, the presence of restaurants and shops and so on. We show that places that offer higher wages also have higher rents, which may entirely offset gains in wages. Self-reported life satisfaction and happiness are similar across the country – places where people have better labour market outcomes are not generally places where people are happier.³ Spatial differences in the trade-off between labour market outcomes, costs of living and amenities help explain these findings on well-being as well as some of the demographic differences observed across the country – young people are concentrated in cities which offer better labour market prospects, whilst older people are concentrated in coastal and rural areas which offer natural amenities and low costs of living.

All of this has important consequences for policy aimed at narrowing economic disparities between areas – 'levelling up' in the language of the current government. We suggest that generating a lot more economic opportunity outside London and the South East will require spatially concentrated investments to offset the self-reinforcing advantages of high-skilled areas. Our analysis suggests that the impact of such investments on spatial disparities in labour market performance will tend to be small, unless they significantly alter the composition of the workforce across areas. Policy may also need to support increased mobility and undertake other complementary investments to help households access the opportunities created by investment.

Places matter to people. Spatial disparities are also important because local social and economic conditions affect individual outcomes. But it is possible to overstate the importance of place in determining economic outcomes and life chances. Spatial disparities also reflect individual inequality. As already noted, the link between individual and spatial disparities is complicated by the fact that people move around. This means that policies that are place-based – that is, targeted at specific places – do not necessarily end up benefiting the people that they aim to help. For many policies, using a place-based approach to targeting policy will also fail to address within-area disparities or reduce individual inequalities which, as shown below, are much bigger than between-area disparities.

The rest of this chapter is structured as follows. Section 2 describes spatial disparities across labour markets and how these disparities have changed over time. Section 3 examines the extent to which spatial disparities in employment and wages reflect differences in the types of people who live and work in different areas. Section 4 considers what drives these differences in composition by looking at both the demand for, and the supply of, different kinds of skills and the

³ There are also important differences in health outcomes and life expectancy across the country (see, for example, Office for National Statistics (2020)). These are beyond the scope of this chapter and will be considered in another chapter of the IFS Deaton Review (Case and Kraftman, forthcoming).

way in which these interact to determine spatial disparities. Section 5 looks at whether spatial disparities in labour market outcomes translate into differences in well-being across areas. The final section discusses policy implications and concludes.

2. Patterns of spatial disparities

This section describes spatial disparities across the country in age, education and economic outcomes. We focus on disparities across travel to work areas (TTWAs). These are constructed to reflect relatively self-contained local labour markets or 'commuting zones' – at least 75% of a TTWA's resident workforce work in that TTWA, and at least 75% of people who work in a TTWA also live in that TTWA. People commute to jobs within TTWAs. To the extent that local conditions drive individual labour market outcomes, these effects would be expected to operate at the level of TTWAs – for example, the relevant labour market for someone living in Salford is Manchester, not Salford. While this justifies focusing on TTWAs, as shown later in this section, there is considerable inequality within TTWAs, and the extent of measured spatial disparities is larger at smaller spatial scales.

We use local authority approximations of TTWAs defined using the 2011 census. This allows us to go further back in time, and link to data sets that are only available at the local authority level. Appendix A details the construction of TTWAs and data sources and Appendix B provides supplementary figures.

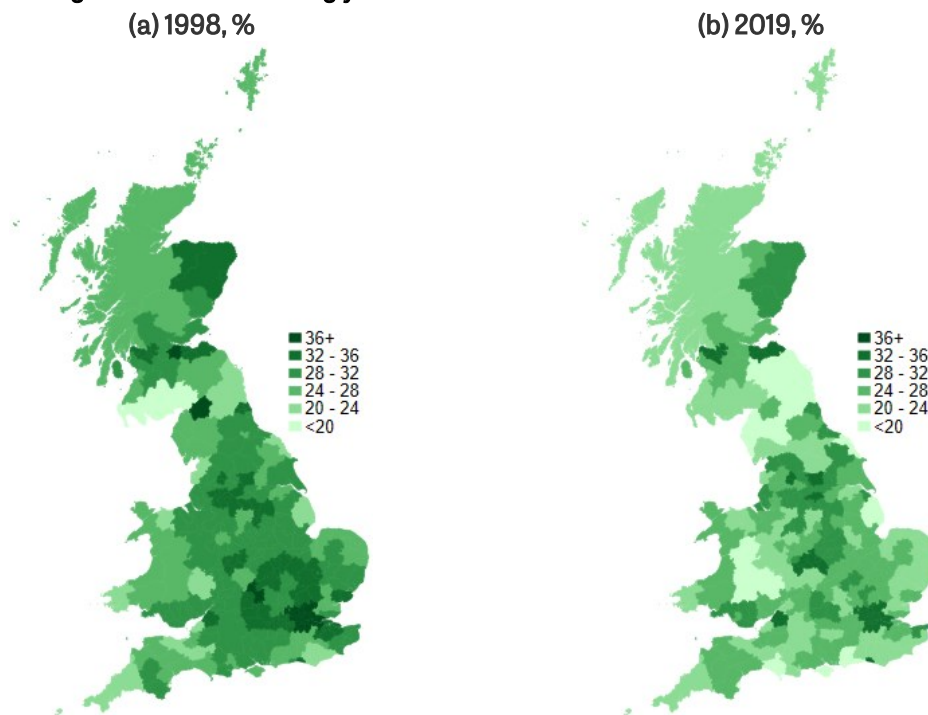
Spatial disparities in age

Labour market outcomes such as wages, employment and hours worked vary over the life cycle. Given lifetime patterns of working, the extent to which labour market outcomes matter for household incomes will thus depend on the age structure of the population.

The age structure of the population differs across the country. Figure 1 shows that young people are disproportionately concentrated in cities. In 2019, 35% of adults living in London and Bristol were under the age of 35, compared with 28% nationally. Other major cities including Manchester, Glasgow, Edinburgh, Leeds and Birmingham also had above-average shares of young adults, ranging from 31% to 34%. At the other extreme, only 17% of adults in Skegness and Louth and in Dorchester and Weymouth were under 35 in 2019. Young adults have become increasingly concentrated in cities over the last 20 years, as can be seen by comparing panels a and b of the figure. Whilst the national share of adults under 35 fell by 3 percentage points between 1998 and 2019, the share of young adults living in London fell by 2 percentage points, and the shares of young adults in Edinburgh, Leeds, Birmingham and Bristol increased.

Figure 2 shows that older people are disproportionately concentrated in coastal and rural areas. In 2019, 23% of adults nationally were aged 65 and over, while in areas such as the Isle of Wight, Torquay and Paignton, and Chichester and Bognor Regis, the share of older adults ranged between 35% and 37%. In contrast, only 16% of adults living in London were 65 or older. Nationally, the share of adults aged 65 or older rose by 3 percentage points between 1998 and 2019. The rise in this share was larger in many coastal and rural areas, whilst the share of older adults in London fell over this period.

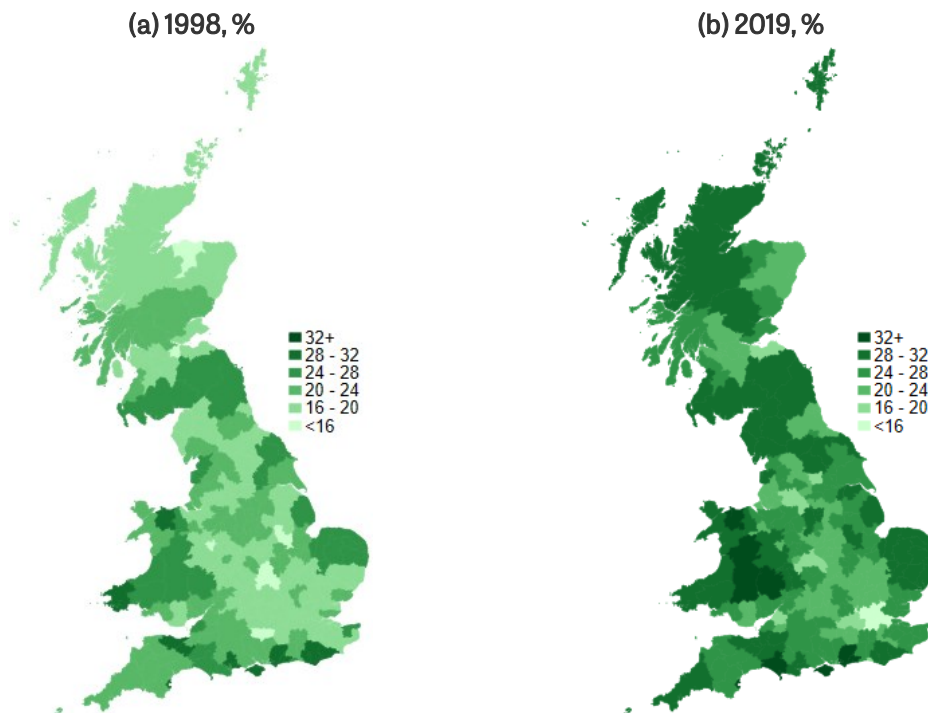
Figure 1. Young adults are increasingly concentrated in certain cities



Note: Adults defined as those aged 16–17 and not in full-time education, and those aged 18 and over. Maps constructed using local authority approximations of 2011 TTWAs as discussed in Appendix A.

Source: Labour Force Survey 1998; Annual Population Survey 2019.

Figure 2. Older adults are concentrated in coastal and rural areas



Note: Adults defined as those aged 16–17 and not in full-time education, and those aged 18 and over. Maps constructed using local authority approximations of 2011 TTWAs as discussed in Appendix A.

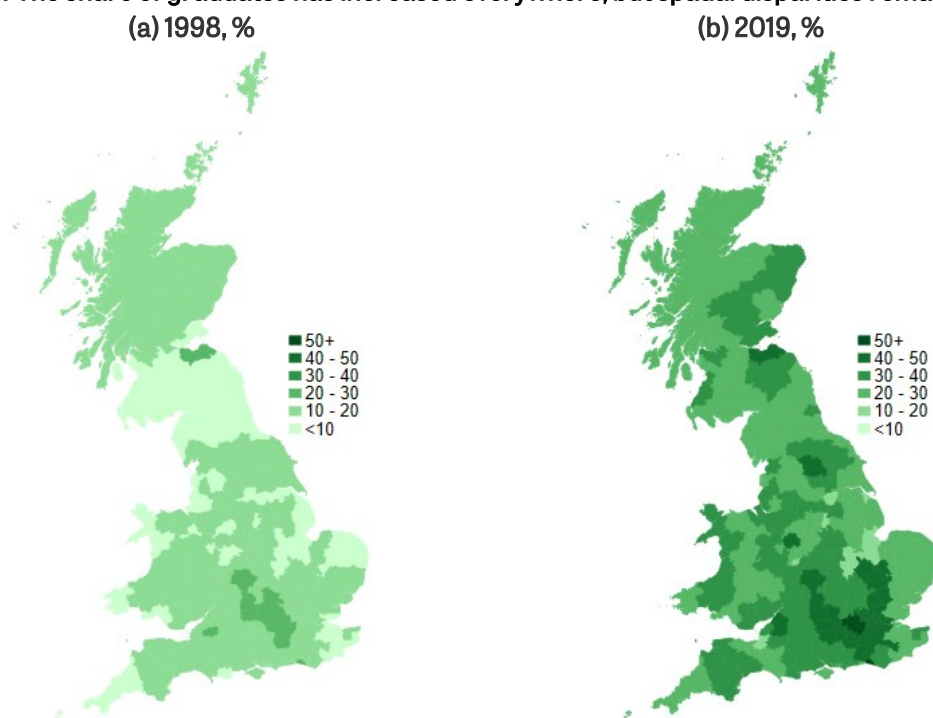
Source: Labour Force Survey 1998; Annual Population Survey 2019.

As discussed later, these spatial disparities in age of the population partly reflect the different trade-offs places offer in terms of labour market opportunities, costs of living and amenities. As we will show, young adults, particularly those with a degree, tend to prioritise labour market prospects when choosing where to live and may also prefer the types of amenities – for example, bars and restaurants – that cities offer. Labour market opportunities matter less for those aged over 65, who are drawn to coastal and rural areas by the natural amenities and low costs of living that they offer.

Spatial disparities in education

As with age, disparities in education matter for labour market outcomes. Figure 3 shows that levels of education also vary substantially across the country. In 2019, graduates made up around half of the working-age population in London and Brighton, compared with less than a fifth of working-age adults (15–18%) in places such as Doncaster, Mansfield and Grimsby. Despite a large expansion of higher education, these spatial disparities in graduate shares have persisted over the last 20 years, as illustrated by panels a and b of the figure. Despite different starting points, areas saw similar percentage point increases in graduate shares between 1998 and 2019.⁴

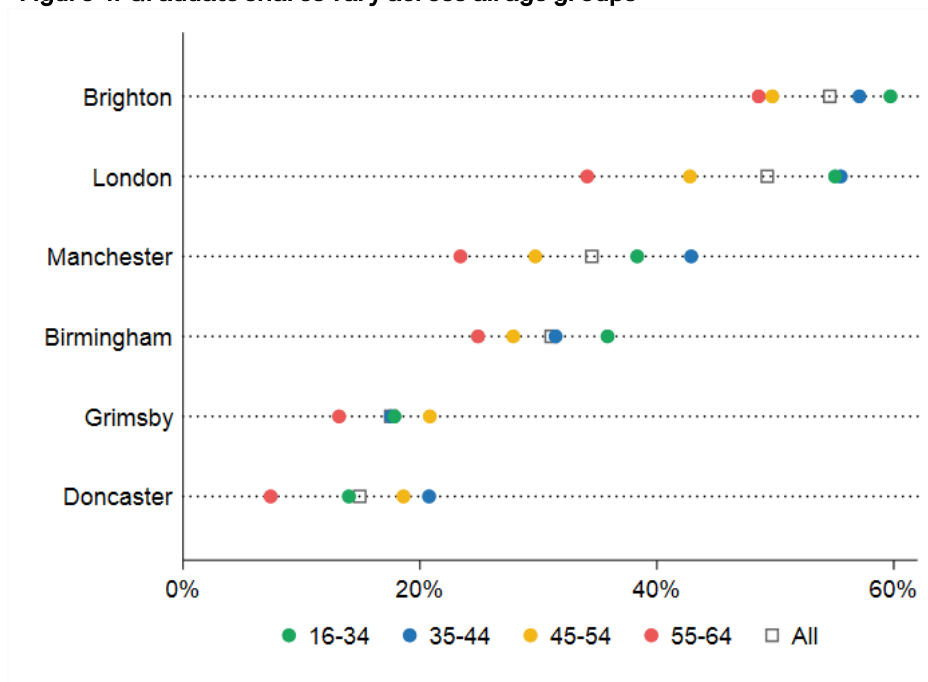
Figure 3. The share of graduates has increased everywhere, but spatial disparities remain



Note: Working-age (16–64) population. Excludes those in full-time education. Maps constructed using local authority approximations of 2011 TTWAs as discussed in Appendix A.

Source: Annual Population Survey.

⁴ This is shown more clearly in Appendix Figure B1, which plots the TTWA-level graduate share in 1998 against the share in 2019. Of course, similar percentage point increases across areas mean that the percentage increase in graduates was larger in areas that initially had lower graduate shares.

Figure 4. Graduate shares vary across all age groups

Note: Figure shows share with degrees in working-age (16–64) population. Excludes those in full-time education.

Source: Annual Population Survey.

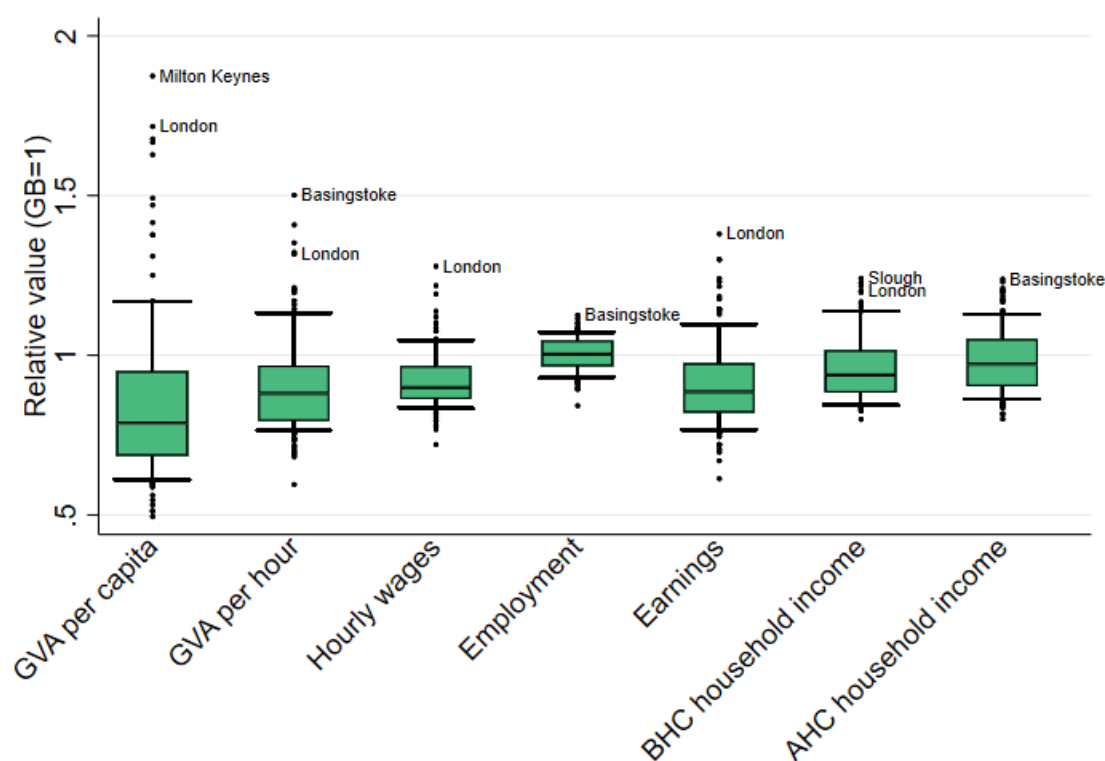
Given that younger people are more likely to have a degree, it is possible that differences in graduate shares could simply reflect differences in the age distribution across areas. Figure 4 shows that this is not the case. Places with low overall graduate shares, such as Doncaster and Grimsby, have lower graduate shares in each age group than places with high overall graduate shares, such as London and Brighton.

As discussed in more detail later, spatial disparities in education reflect differences in educational attainment of children who grow up in different areas and selective patterns of mobility after graduation. Differences in educational outcomes across areas, and graduates' choices on where to live and work, are both a cause and a consequence of differences in labour market outcomes across the country.

Spatial disparities in productivity, earnings and incomes

Figure 5 shows the extent of spatial disparities in different economic outcomes across TTWAs in 2018, the latest year for which data on all outcomes are available. For each outcome, the 'box and whisker' plot summarises the distribution of area averages (means). The upper and lower bounds of the 'box' plot the 25th and 75th percentiles of area averages for that outcome, and the line in the middle of the box plots the median. The lower and upper lines on the 'whiskers' on either side of the box correspond to the 10th and 90th percentiles, respectively. Areas in the top or bottom 10% of area averages are plotted with a dot. Outcomes are shown relative to the relevant national average for Great Britain. For example, the average hourly wage in London is £21 and the national average wage is £16, so London has a relative value of 1.3 for hourly wages.

Figure 5. The extent of spatial disparities and the ranking of individual areas depends on the outcome measure



Note: Data for 2018. Figures constructed using local authority approximations of 2011 TTWAs as discussed in Appendix A. The highest-value TTWA and London are labelled. Slough refers to the 'Slough and Heathrow' TTWA. Earnings are imputed as TTWA-level employment rate (APS) times TTWA-level average employee earnings (ASHE). We do not have reliable data on self-employed earnings by TTWA. Hourly wages, employment and weekly earnings are defined for working-age (16–64) population. Household incomes are equalised to reflect differences in household size and composition. Estimates for household incomes are for England and Wales only, but this is not why area averages are more compressed than for earnings, as shown in Appendix Figure B2 which reproduces Figure 5 for England and Wales only. GVA is gross value added. BHC stands for before housing costs and AHC for after housing costs.

Source: Office for National Statistics (GVA per capita, GVA per hour and household incomes); Annual Population Survey, APS (employment and earnings); Annual Survey of Hours and Earnings, ASHE (hourly wages and earnings).

The first box and whisker plot shows the distribution of gross value added (GVA) per capita across TTWAs. This measures the value of economic output produced in a TTWA, divided by the population of that TTWA. The plot shows quite wide dispersion and a considerable skew at the upper end, as evidenced by the median being far below the mean. Milton Keynes, the area that is top on this measure, has GVA per capita nearly double (88% higher than) the national average. GVA per capita in London, the second-highest area, is 72% higher than the national average. At the other end of the spectrum, Kilmarnock and Irvine, Hastings, and Torquay and Paignton have GVA per capita that is half the national average.

GVA per capita, or the closely related gross domestic product (GDP) per capita, is one of the most widely used measures of local economic performance. It is often used in international comparisons of spatial disparities (see Box 1). However, it can be a poor measure of local productivity – because the output *produced* in an area is divided by the population, it can be severely distorted by commuting flows. Using relatively self-contained areas, as we do, helps

address this problem, but does not eliminate it – GVA per capita in Milton Keynes and London is partly high because of high levels of commuting into these areas. Further, employment and hours worked vary across areas, partly reflecting the differences in demographics discussed above. The areas with the lowest GVA per capita in Figure 5 illustrate this nicely. Hastings and Torquay & Paignton are home to lots of older people, many of whom will be retired: 31–33% of their adult populations were aged 65 or older in 2018, compared with just 23% nationally.

GVA can be used to get a measure of differences in labour productivity provided that it is normalised by hours worked, rather than population, to allow for differences in commuting flows, employment rates and hours worked across areas. The second plot in Figure 5 shows the extent of spatial disparities in labour productivity as measured by GVA per hour. The dispersion continues to be quite wide, though less so than for GVA per capita. The skew is much less pronounced at the upper end, suggesting that GVA per capita figures are distorted by commuting patterns, even at the level of TTWAs (there are also some minor differences in employment and hours worked). Using GVA per hour worked rather than per capita sees Milton Keynes drop from the highest-ranked TTWA in Britain to seventh, and London from the second-highest to fifth.

However, while GVA per hour is a better measure of labour productivity than GVA per capita, differences in productivity do not necessarily translate into differences in wages. GVA captures the total value of output produced in an area, which must be used to reward all factors of production, not just labour. From this output, profits must be paid to shareholders, rents to landowners and interest to capital owners. If the share of GVA that is paid to employees in the form of wages varies across areas, or if ownership of assets is distributed differently across areas – for example, if profits from Aberdeen oilfields go to shareholders in London – then GVA per hour will be a poor proxy for local wages.

The third plot shows that there is much less variation in average hourly wages than in GVA per hour. London comes out top, with average wages around a third higher than the national average. Average wages in the lowest-wage area, Scarborough, are around a quarter lower than the national average. These are still sizeable differences, but far less stark than differences in GVA per capita or GVA per hour.⁵

Employment rates also vary across TTWAs, although differences are relatively small. This can be seen in the fourth plot in Figure 5. Taken together, spatial variation in wages and employment (the probability of having a job), as well as in hours worked (Schlüter, 2013), leads to spatial variation in earnings. The distribution of average earnings is summarised in the fifth plot. London has the highest weekly earnings – suggesting, unsurprisingly, that differences in earnings correspond better with the common view of spatial disparities than the more abstract productivity differences captured by GVA per capita or per hour worked. As wages and employment rates (as well as hours) are positively correlated, differences in earnings are larger than those for either component individually. Differences in earnings largely come from differences in wages and employment rates, so these are the outcomes focused on in the rest of the chapter.

⁵ Our wage data (from the Annual Survey of Hours and Earnings) only cover employees. If self-employed wages are more spatially dispersed than employee wages, Figure 5 may understate the extent of spatial wage disparities. Appendix Figure B3 suggests we might expect more spatial variation in wages if self-employed workers were included in the data, since there is more variation in the skills of self-employed workers (whether they have degrees) than in the skills of employees across areas.

Box 1. The most spatially unequal country in the developed world?

It is often said that the UK is the most spatially unequal country in the developed world. This is repeated so often in the media and public discourse that it is widely accepted as fact. However, comparing spatial disparities across countries is difficult, and commonly used measures are particularly problematic for the UK.

The studies that find the largest spatial disparities in the UK, relative to other countries, compare differences in GDP per capita across small administrative areas – so-called territorial level 3 or TL3 regions (McCann, 2019; Carrascal-Incera et al., 2020; Davenport and Zaranko, 2020). Countries are ranked by the difference in GDP per capita in the top and bottom TL3 regions, or by some dispersion measure calculated using the whole distribution. There are two major problems with these comparisons:

- The size of TL3 regions varies widely depending on how administrative boundaries are drawn. The UK has 179 TL3 regions, the second largest number of all OECD countries, in contrast to France which has 96 TL3 regions and Spain which has 59. Most importantly, the UK is unique in that its most productive city, London, is split into more than 20 separate TL3 regions. This means that in the UK, the top regions consist solely of cities – often just different parts of London – whilst regions at the bottom consist solely of rural areas. In countries with fewer TL3 regions, cities are often grouped with their suburbs and surrounding rural areas, so differences between regions are muted.
- As discussed in the main text, GDP per capita is a flawed measure of productivity (let alone living standards) because it divides the output produced in an area by the number of residents who live in the area. This is exacerbated when looking at small areas, such as the separate TL3 regions in London. Camden & City of London – the highest-ranked TL3 region – has a population of around 260,000, yet some 800,000 people work there and contribute to its GDP (Selby-Boothroyd, 2018). The use of GDP per capita, combined with the artificial division of London into 21 separate areas, vastly overstates the level of spatial disparities in the UK.

McCann (2019) compares the UK and other OECD countries using several different spatial levels. In the comparisons that use GDP per capita at the TL3 level, the UK consistently comes out top in terms of spatial disparities. In contrast, comparisons at the TL2 level – a larger level of aggregation where London is counted as a single region, comparable to Paris, Berlin or Tokyo – put the UK in the top quarter or top fifth of countries. Of course, the administrative boundaries of London still leave out large numbers of commuters from surrounding areas. Comparisons that use functional labour market areas – 'metro urban regions', similar to TTWAs – place the UK around the middle of the pack.

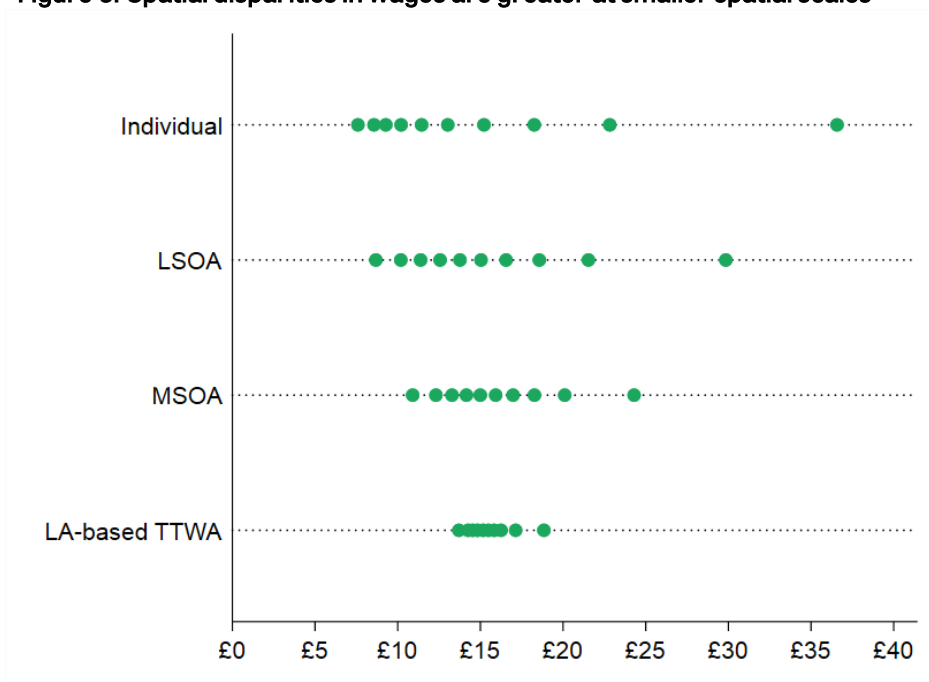
Taken together, the evidence suggests that spatial disparities in the UK are relatively high by international standards, but nowhere near as high as sensationalist headlines might have us believe.

Individual inequalities and inequalities at smaller spatial scales

This chapter focuses on differences across local labour markets – as these are, by definition, the most relevant spatial scale at which to consider disparities in labour market outcomes. As discussed above, to the extent that local conditions affect individual outcomes in employment and wages, this would be expected to happen at the level of TTWAs, rather than at smaller spatial scales. As a result, differences between neighbourhoods within a TTWA – between, say, Peckham and Dulwich (two neighbouring areas in South-East London) – are much more likely to reflect differences in the composition of people than area-specific effects on labour market outcomes.

Whilst focusing on TTWA makes sense from an analytical perspective, it masks considerable inequality in outcomes across individuals and at smaller spatial scales. This is illustrated by Figure 6, which plots average wages in 2019 across deciles of individuals and areas at increasing levels of aggregation: Lower Layer Super Output Areas (LSOAs) of around 500–1,000 households each, Middle Layer Super Output Areas (MSOAs) of around 2,000–6,000 households each, and our 136 local-authority-based TTWAs.

Figure 6. Spatial disparities in wages are greater at smaller spatial scales



Note: Shows average wages among working-age (16–64) employees in 2019 by decile of individuals and areas, where areas are ranked by their mean wage. Areas are based on place of residence.

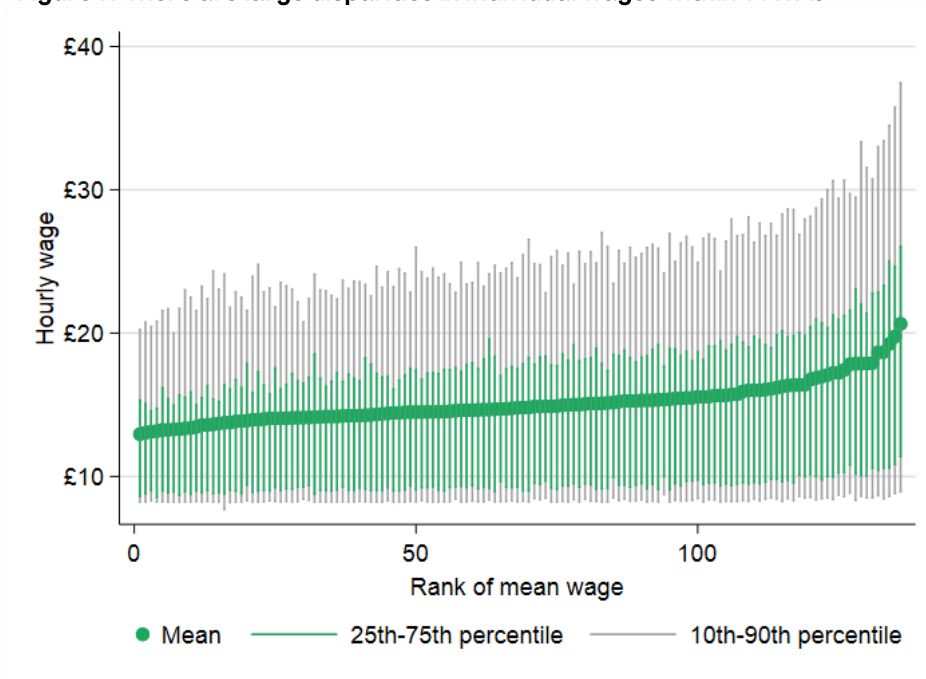
Source: Annual Survey of Hours and Earnings.

The figure shows that differences between the deciles of the wage distribution are largest at the individual level and fall consistently when moving to higher spatial scales, reflecting the fact that mobility – and hence the extent of segregation by income – falls as spatial scale increases (Manning and Petrongolo, 2017).

Focusing on explaining disparities across TTWAs also masks the extent of variation within local areas, as shown in Figure 7, which ranks TTWAs by average wage in 2019 and plots the distribution of wages within each TTWA. Within-area disparities are much bigger than between-area disparities. Figure 7 also highlights that differences in average wages across areas are

driven more by the top of the wage distribution – the wages of the top 25% and top 10% rise steeply moving across areas. Variation at the very top of the wage distribution is greater still. The super-rich are highly concentrated in London, which is home to 35% of the top 1% of income taxpayers but only 10% of the population (Joyce, Pope and Roantree, 2019). In contrast, there is little variation in wages at the bottom of the wage distribution. Wages of the bottom 10% are essentially the same across areas – unsurprisingly given the bite of the minimum wage – and the same is true for wages of the bottom 25%.⁶ We return to these points, and their implications for policy, below.

Figure 7. There are large disparities in individual wages within TTWAs



Note: Shows distribution of wages among working-age (16–64) employees across TTWAs in 2019, ranked by TTWA-level mean wage.

Source: Annual Survey of Hours and Earnings.

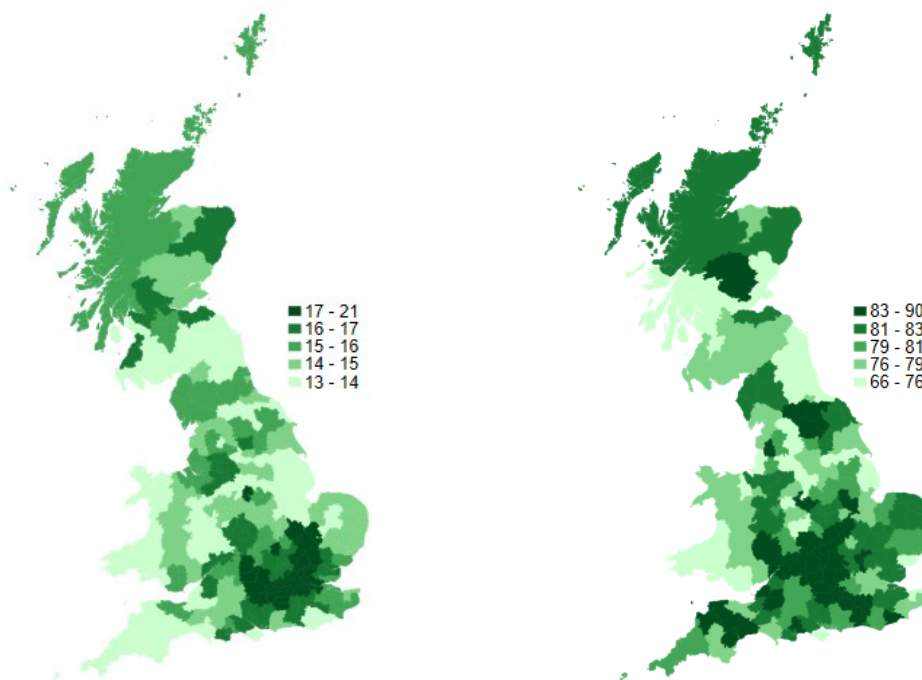
Spatial disparities in wages and employment over time

Having described spatial disparities across different outcomes, we now turn to our main measures of labour market outcomes – wages and employment. Figure 5 summarised the dispersion of wages and employment rates across TTWAs. Figure 8 illustrates these disparities in more detail by mapping outcomes across the country. In 2019, the average wage in London was over £20, whereas the average wage in Scarborough was just £13. Employment rates ranged from 66% in Skegness and Louth to 90% in Harrogate.

The North–South divide is visible, as is the relatively poor performance of many coastal and some rural areas. Differences in wages are more pronounced than differences in employment. Areas that have lower wages do not always have lower employment rates, but some areas – such as the South of Scotland, North of England, Lincolnshire and Wales – fare poorly on both measures.

⁶ This implies that wage inequality is higher in TTWAs with higher average wages. Using administrative payroll and benefits data, Rae and Nyanzu (2019) find that household incomes also vary greatly within TTWAs, with more household income inequality in richer TTWAs.

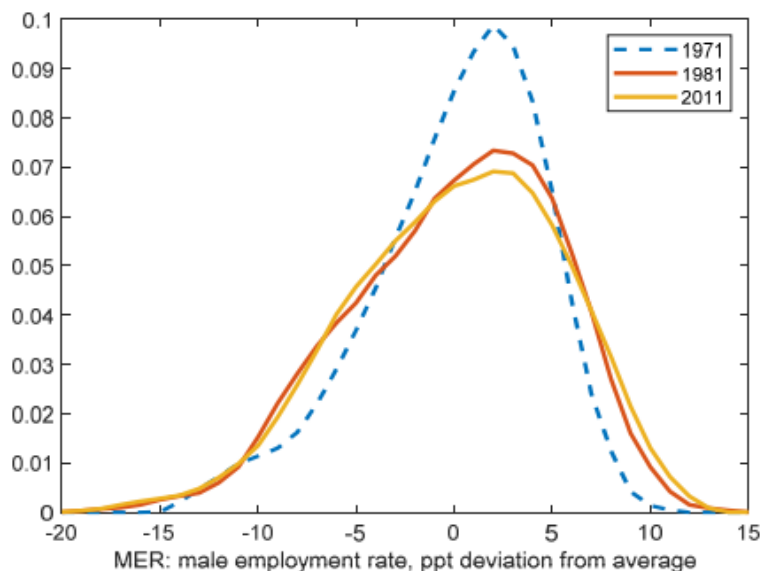
Figure 8. Average wages and employment rates vary significantly across areas
 (a) Wages, £ per hour (2019) (b) Employment, % (2019)



Note: Maps constructed using local authority approximations of 2011 TTWAs as discussed in Appendix A. Working-age (16–64) population. Wage data exclude self-employed.

Source: Annual Population Survey; Annual Survey of Hours and Earnings.

Figure 9. Spatial disparities in male employment increased in the 1970s



Note: Shows male employment rates at the local authority level.

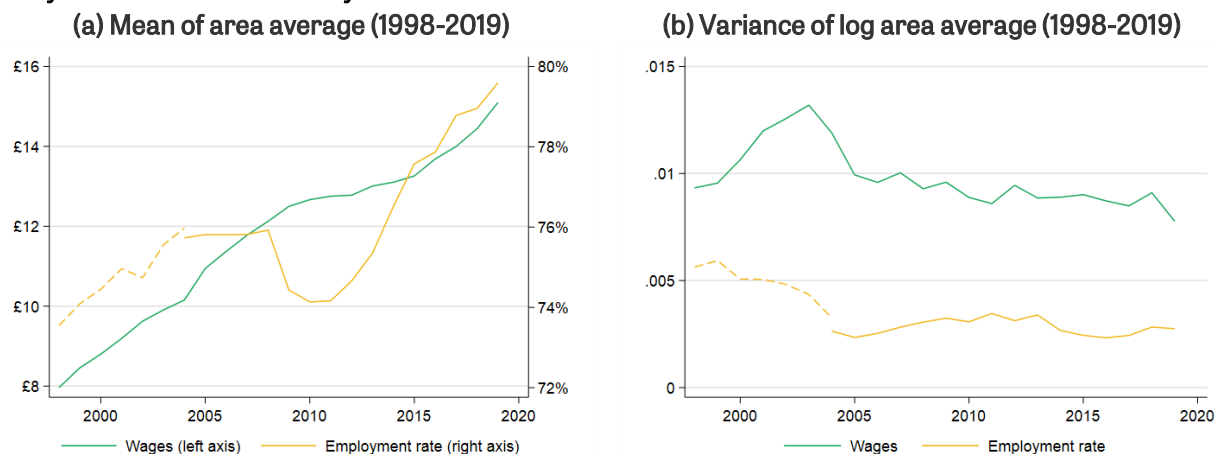
Source: Rice and Venables, 2021.

How have these patterns changed over time? We focus on the period from 1998, when we have good data on individual wages and employment that help us better understand the nature of spatial disparities and why they persist. These data paint a rich picture of the causes and consequences of these spatial disparities for individual workers.

Starting in the late 1990s, however, misses a key part in the story of spatial disparities in the UK played by the dramatic fall in manufacturing employment in the 1970s. Employment in the secondary sector (manufacturing, construction and utilities) fell from its peak of 40% in 1996 to 30% in 1981 and continued to decline by 4–5 percentage points per decade to reach 15% by 2015. This shock was highly spatially concentrated – for example, some areas saw their male employment rates fall by 5–10 percentage points between 1971 and 1981. Figure 9, from Rice and Venables (2021), plots the distribution of male employment rates by local authority around the national average for three different census years. It shows that spatial disparities in male employment rates increased between 1971 and 1981, and that this increase persisted up to 2011. Appendix Figure B4, also from Rice and Venables (2021), shows that areas that were badly affected by deindustrialisation were still feeling the effects by 2011.

Data on GDP per capita – a similar measure to the GVA per capita used above – also reflect the spatial effects of deindustrialisation. Appendix Figure B5 shows that GDP per capita in the West Midlands – a region heavily reliant on manufacturing – went from 3% above the UK average in 1971 to 9% below in 1981. By 1996, it was still 7% below the national average. The spatially uneven and persistent effects of deindustrialisation in the 1970s provide the backdrop to our period. Although we do not consider this period directly, the evidence provided on the spatial concentration of skills and the self-reinforcing nature of spatial differences in the demand for and supply of skills helps explain why the effects of 1970s deindustrialisation have persisted.

Figure 10. Spatial disparities in wages and employment rates have fallen and are around as low as they have been in the last 20 years



Note: Working-age (16–64) population. Note that from 2004 onwards, the Labour Force Survey (LFS) was supplemented by a boost sample to make it representative at the local-authority level, forming the Annual Population Survey (APS). Employment rates calculated using the LFS are shown by the dotted yellow line and those using the APS by the solid yellow line. Notice that the expansion of the sample size in 2004 artificially reduced the variance of log TTWA averages in panel b. However, as the dotted line shows, the fall in measured disparities in 2004 does not simply reflect this change.

Source: Labour Force Survey (pre-2004 employment); Annual Population Survey; Annual Survey of Hours and Earnings.

From 1998 onwards, we can use microdata to look at changes to local wages and employment rates. Figure 10 shows the mean and variance across TTWAs over time. Panel a shows that the average nominal wage increased throughout the period, though real wages (adjusting for inflation) fell after 2008 and then stagnated (Giupponi and Machin, forthcoming). Average employment rates also increased for most of the period, with a temporary fall during the financial crisis.

Panel b shows the extent of spatial disparities, measured by the variance of log TTWA averages, which is invariant to common growth in wages across areas. Spatial differences in wages increased in the early part of the 2000s before falling back just before the financial crisis. They have been on a slow downward trend since. Two factors are likely to play a role in this convergence: a fall in wages for those at the top of the distribution since the financial crisis, and increases in the minimum wage – in particular, since 2016 – which have pushed up wages for those at the bottom of the distribution (see Agrawal and Phillips (2020) for evidence on both). Since the former group account for a higher share of employment in areas with high average wages, and the latter for a higher share in areas with low average wages, these factors serve to narrow spatial disparities in average wages. The expansion of higher education, which resulted in larger percentage increases in graduate shares in areas with initially lower shares of graduates – see Appendix Figure B1 – may also play a role.

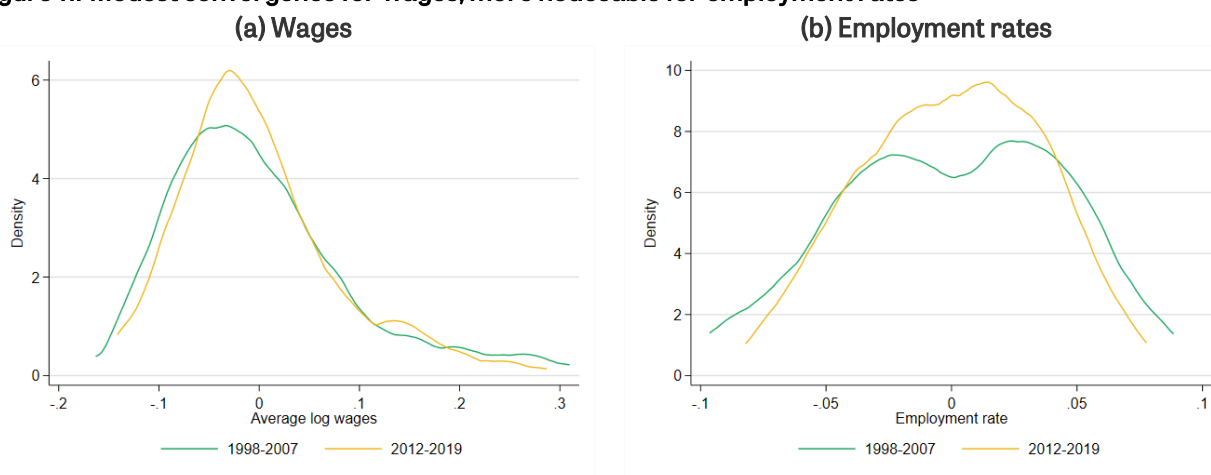
In contrast to wages, where the modest fall in spatial disparities began in the early 2000s, disparities in employment rates fell markedly in the early 2000s but have been broadly stable since. Areas with low employment rates in 1998 saw large increases in employment between 1998 and 2004, as shown in Appendix Figure B6. During this time, the New Labour government introduced several policies that boosted employment rates, in particular the working families' tax credit in 1999 and the working tax credit in 2003 (Blundell et al., 2000; Mulheirn and Pisani, 2008). These policies appear to have benefited low-employment areas the most, contributing to a convergence in employment rates.

Figure 10 shows that spatial disparities in wages and employment are around as low as they have been in the last 20 years. This overall picture of a gradual reduction in disparities could hide bigger changes for the best- and worst-performing areas. Figure 11, which plots the entire distribution of area wages and employment rates for periods before and after the Great Recession, shows that this is not the case. These overall distributions repeat the pattern for variances: a slight tightening of the spatial distribution of wages and a more marked tightening for employment rates.

While one might argue that this improvement in overall spatial disparities – modest for wages, more noticeable for employment rates – runs counter to public perception, there is another important aspect to consider. Where do areas fit within the overall distribution? Here, the story is one of considerable persistence over time. Most areas that were struggling 20 years ago with relatively low wages and employment rates are still struggling today, and vice versa for areas that were flourishing.

This is shown in Figure 12, which plots area averages before the Great Recession against area averages after the Great Recession. Average wages across areas are highly persistent, with wages in 1998–2007 explaining 96% of the variation in wages in 2012–19. Employment rates are also persistent, though consistent with what has been seen so far, somewhat less so than wages.

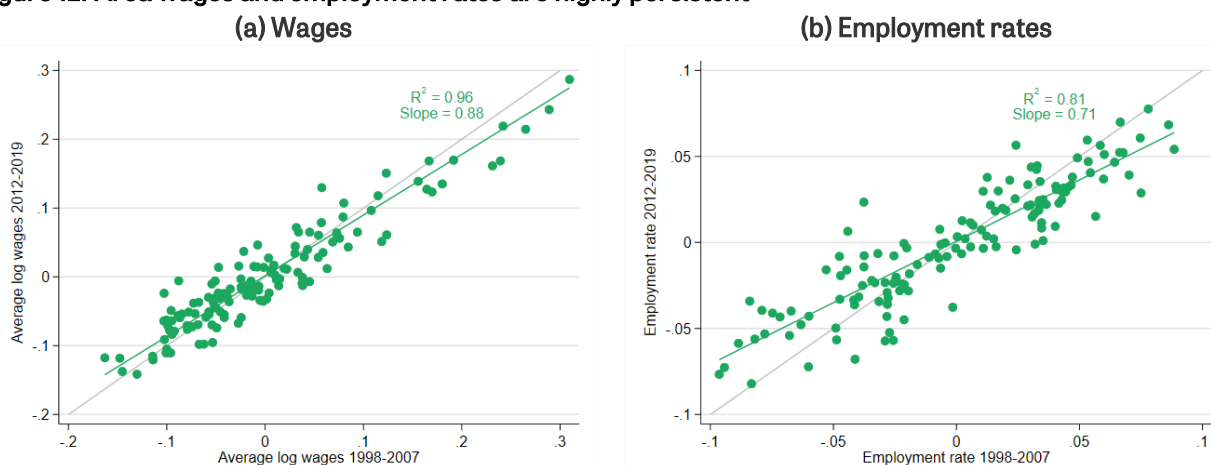
Figure 11. Modest convergence for wages, more noticeable for employment rates



Note: Data are pooled across two periods – before the Great Recession (1998–2007) and after the Great Recession (2012–19) – normalising outcomes so that the mean across areas is zero in each period. Working-age (16–64) population. Dropping data for the financial crisis and pooling over time smooths out temporary local shocks, to focus on longer-run structural changes. The two peaks in the distribution of employment rates result from the overall growth in employment rates between 1998 and 2007.

Source: Labour Force Survey (pre-2004 employment); Annual Population Survey; Annual Survey of Hours and Earnings.

Figure 12. Area wages and employment rates are highly persistent



Note: Working-age (16–64) population. As for Figure 11, data are pooled across two periods – before the Great Recession (1998–2007) and after the Great Recession (2012–19) – normalising outcomes so that the mean across areas is zero in each period.

Source: Labour Force Survey (pre-2004 employment); Annual Population Survey; Annual Survey of Hours and Earnings.

What explains spatial disparities in the labour market?

The structural shift from manufacturing to services has had a profound effect on the economic geography of the UK. The effects of this shift have been highly persistent and continue to shape spatial disparities today, despite recent convergence in labour market outcomes – which in part may reflect the impact of the financial crisis and policy changes. But recent trends do not simply reflect policy changes and adaptation to deindustrialisation and the financial crisis – other structural shifts play an important role, often in ways that reinforce spatial disparities. These shifts are discussed further in Section 4.

The way in which individual firms and workers respond to these changes can help explain why spatial disparities in the labour market are persistent and allow us to better understand the causes and consequences of these disparities. To see how, we proceed in three steps.

First, Section 3 uses microdata on individual workers to consider the extent to which spatial disparities reflect differences in the types of people who live and work in different areas. We show that disparities largely reflect the concentration of workers who would have better labour market outcomes wherever they work. Where these high-skilled workers, and the firms that employ them, are concentrated also matters because some areas generate higher productivity and better labour market outcomes. These 'area effects' arise because of things inherent to the place – such as the climate or physical geography – or because of productivity benefits arising from the concentration of firms and workers. We show that skilled workers generally work in better-performing labour markets and vice versa. The concentration of skilled workers in certain areas takes individual labour market advantages and magnifies them further.

Second, Section 4 considers what drives the concentration of skilled workers by looking at both the demand for, and the supply of, different kinds of skills and the way in which these interact to determine spatial disparities. One key driver is the geography of high-skilled jobs. We show that the geography of demand for high-skilled workers is highly spatially concentrated and possibly becoming more so. Although we cannot differentiate between them, the urban economics literature offers various theories for what might be driving these patterns (see Ottaviano and Thisse (2004) and Redding and Turner (2015) for surveys). While the concentration of skills is partly demand-led, spatial disparities in the supply of skilled workers also matter. We provide evidence that disparities in educational attainment and the selective migration of graduates both play a role in explaining these disparities. The demand for, and supply of, skills interact in ways that can be self-reinforcing and hence magnify the effect of shocks and make spatial disparities persistent. We show that this self-reinforcement is particularly pronounced in the highest-wage areas and at the upper end of the wage distribution. The result is that the highest-paid jobs are concentrated in London and a handful of other areas.

Third, Section 5 looks at whether spatial disparities in labour market outcomes translate into disparities in well-being. If people are sufficiently mobile, the utility individuals derive from living in an area – taking all these factors into account – should be broadly equalised across areas (Rosen, 1974; Roback, 1982). Although individual utility is not directly observable, we can look at differences in self-reported well-being across areas. We show that places where people have better labour market outcomes are not generally places where people are happier. This is an important point to bear in mind: labour market differences do not necessarily translate into differences in well-being, or even in living standards (because of differences in the cost of living).

3. Understanding spatial disparities in employment and wages: the role of individuals and areas

The previous section showed that there are large and persistent disparities in labour market outcomes across areas of the UK. This section begins to consider the role that firms and workers play in understanding these disparities by asking whether they reflect differences in the people who live and work in these areas, or whether they reflect area-generated differences in outcomes for the same types of people.

There are two caveats to this approach. First, it focuses on area effects in the labour market, after people have acquired much of their education. The analysis below does not rule out the possibility of area effects on young people's educational attainment – for example, through peer effects, school quality or local incentives to invest in education. That said, while there are large spatial disparities in educational outcomes, and we return to these below, it is important to be cautious in attributing much of a role to place in determining these disparities. To our knowledge, there are no studies that quantify the extent to which differences in education across local authorities or TTWAs reflect differences in school quality and other attributes of place,⁷ while the UK literature considering effects at neighbourhood level is inconclusive on whether such effects occur. In contrast, a large literature highlights the importance of differences in the characteristics of parents.⁸

Second, asking the question this way ignores questions about why these differences emerge. To what extent does the concentration of workers across areas reflect spatial differences in the demand for different types of workers by firms? And to what extent do the location choices of firms reflect spatial differences in the supply of different types of workers? Setting aside these questions for now – we return to them in Section 4 – breaking down the disparities into these two components is still helpful, as illustrated by thinking about two extreme scenarios.

If disparities in wages purely reflect differences in the type of people who work in different areas, then someone moving jobs from low-wage Hull to high-wage London would see no change in their wages. Likewise, investment to replicate London's infrastructure and business environment in Hull would not improve wages, unless it also changed the types of people who worked there. At the other extreme, if disparities purely reflect wage premiums that are specific to the local economy, then a person moving jobs from low-wage Hull to high-wage London would see their wages increase by 44% (the difference in average wages between Hull and London in 2019). Likewise, replicating London's infrastructure and business environment in Hull would increase average wages there by 44%, even if there was no change in the types of people who work there. The two extreme scenarios for employment rates can be thought about in the same way, though of course the percentage gains would be different.

Estimating individual and area effects and understanding their contributions to spatial disparities

The breakdown for observed spatial disparities lies somewhere between these two extremes. In this section, we use regression analysis to decompose individual wages into the parts that can be attributed to individual characteristics – 'individual effects' – and to 'premiums' that are specific to the local economy in which they work – 'area effects'.⁹ Such regressions are commonly used in

⁷ Studies on the determinants of educational outcomes often include local authority (LA) or region dummies as explanatory variables. However, given that existing studies do not set out to estimate area effects in education, they tend to include only rough proxies for parental background (such as eligibility for free school meals) and often control for other variables that should be thought of as part of the area effect (such as per-pupil spending, school type or LSOA-level deprivation measures). Common measures of school performance are published at the LA level; however, these also capture the effect of student composition, because parental background is likely to affect the rate at which children progress as well as their attainment at any point in time (Andrews, 2017).

⁸ The existence and importance of peer effects at the neighbourhood level in the UK are contested (Oreopoulos, 2006; Gibbons, Silva and Weinhardt, 2013; Weinhardt, 2014). Evidence of neighbourhood effects on education from the US (Chetty, Hendren and Katz, 2016; Chetty and Hendren, 2018; Chyn and Katz, 2021) may not be applicable to the UK, given the many structural differences – for example, in the funding of state schools and the generosity of the benefits system.

⁹ Note that we are focusing on area effects in the labour market, *after* people have acquired education and skills. The analysis below does not rule out the possibility of area effects on young people's educational attainment – for example, through peer effects, school quality or local incentives to invest in education.

the urban economics literature – see Combes and Gobillon (2015) for a review and Card, Rothstein and Yi (2021) for a recent application to the US.

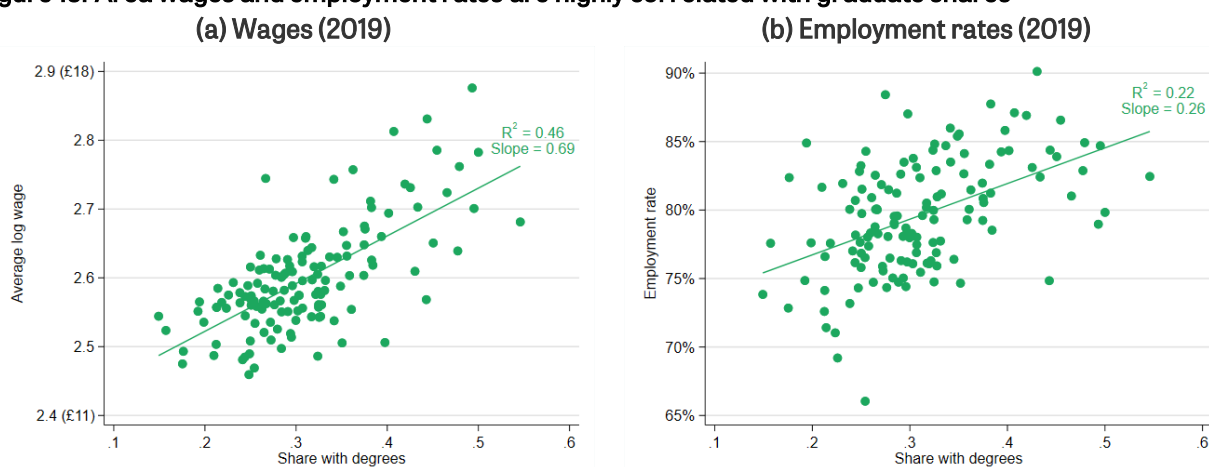
We then use a variance decomposition proposed by Gibbons, Overman and Pelkonen (2014) to ask how much of the observed spatial disparities in wages and employment can be attributed to area effects, to differences in the composition of workers across areas (in terms of their individual effects) and to the correlation between the two.¹⁰ In principle, the correlation between individual effects and area effects could be positive or negative. In practice, the correlation is positive – workers with higher labour market potential wherever they live work in areas with higher area effects, and vice versa.

To illustrate the key ideas, we begin with an example using spatial differences in educational attainment. However, wages and employment outcomes depend on many individual characteristics and our regression analysis considers not just degree status, but a range of individual-level characteristics, including age, gender and other fixed attributes (such as 'ability') that are not observed in the data. We describe the approach non-technically – the estimation equations, variance decompositions and details on implementation are set out in Appendix C.

An example: educational attainment

Better-educated people have higher wages and are more likely to be employed (Blundell et al., 2018). Educational attainment of the working-age population also varies significantly across the country, as illustrated in Figure 3. These two facts suggest that spatial disparities in labour market outcomes may partly reflect where educated people live and work.

Figure 13. Area wages and employment rates are highly correlated with graduate shares



Note: Working-age (16–64) population.

Source: Annual Population Survey; Annual Survey of Hours and Earnings.

¹⁰ The academic literature refers to these differences in composition as arising from the 'sorting' of different types of workers into different places. Such sorting is an equilibrium concept and can occur because of the intergenerational transmission of individual characteristics, the effect of place on individual characteristics (e.g. differences in school quality) or selective mobility across places based on individual characteristics. Unfortunately, it appears the term can lead to confusion in public debate if it is taken to only refer to selective mobility between places.

Consistent with this, Figure 13 shows that TTWA-level wages and employment rates are highly correlated with the share of the working-age population with a degree. Nearly half of the variation in average wages across areas is accounted for by variation in this coarse measure of education. Degree shares matter less for employment, where other factors such as ethnicity and age play a larger role, though they still account for nearly a quarter of the variation in employment rates across areas.

How big are spatial differences in wages and employment allowing for differences across areas in individual characteristics?

We start by comparing 'raw' differences in average wages between TTWAs, pooling across the post-Great-Recession years (2012–19). The first row of numbers in Table 1 summarises the distribution of these raw area differences. Each column reports the percentage change in wages when comparing areas at different parts of the distribution of estimated area effects (Appendix Figure C1 plots the entire distribution).

Table 1. Spatial disparities in wages and employment rates look smaller when moving away from comparisons at the extremes of the distribution and when allowing for differences across areas in individual characteristics

	Max–Min	Med–Min	Max–Med	P90–P10	P75–P25
Wages (log × 100, 2012–19)					
Raw	43	13	30	21	9
Controlling for observable characteristics (Mincerian)	28	8	20	13	5
Controlling for time-fixed observable and unobservable characteristics (AKM)	19	9	10	6	3
Controlling for time-fixed characteristics and time-varying observables (AKM with controls)	17	8	9	6	3
Employment rates (ppt, 2012–19)					
Raw	16	8	8	10	6
Controlling for observable characteristics (Mincerian)	11	6	5	6	4

Note: Working-age (16–64) population. Details of the underlying regression specifications are given in Appendix C. Results for wages show the difference in log wages between areas, which is approximately equal to the percentage difference in wages. Results for employment show the percentage point difference in employment rates between areas.

Source: Annual Population Survey; Annual Survey of Hours and Earnings.

Differences in average wages across areas are quite large, particularly comparing the extremes of the distribution. Wages in the highest-wage area (London) are 43% higher than in the lowest-wage area (Thetford and Mildenhall).¹¹ The second and third columns show that the distribution is skewed: the median area has average wages only 13% higher than the lowest-wage area, but wages in the highest-wage area are 30% higher than for the median area. The final two columns show that comparing the extremes is misleading. The area at the 90th percentile has wages that

¹¹ This is based on average wages over the 2012–19 period. Comparisons earlier in the text between London and Scarborough are based on 2019 alone.

are only 21% higher than the area at the 10th percentile. Wages in the area at the 75th percentile are only 9% higher than wages in the area at the 25th percentile.

The next three rows of Table 1 make the same comparisons, but using estimates of the area effect on wages controlling for an increasingly rich set of worker characteristics. Details of the underlying regression specifications are in Appendix C.

The second row compares estimated area effects controlling for gender, age, skill level and full-time/part-time work status.¹² We refer to these estimates as 'Mincerian'.¹³ Controlling for these differences reduces the differential for each of the five comparisons by roughly 40% – a lot of the differences in wages across areas reflects differences across areas in this relatively small set of characteristics.

An intuitive way to interpret these 'Mincerian' area effects is as estimating what would happen when someone moves jobs across different areas holding constant their gender, age and so on. As our wage data are a panel which tracks workers over time, we do not need to rely on this thought experiment. Instead, we can look directly at what happens to the wages of workers as they move jobs around the country. Area effects estimated by following workers over time – which we refer to as 'AKM'¹⁴ – capture the average change in wages experienced by a worker when they move to London, or to Hull, or to any of our 136 TTWAs.

The advantage of these AKM estimates is that they hold constant everything about workers that is fixed across time whether it is recorded in the data (observable) or not (unobservable). For example, while the Mincerian estimates hold constant gender – observed in the data – the AKM estimates also hold constant ethnicity and ability – unobserved in the data. We can further control for observed characteristics that change over time in a way that may correlate with job moves between areas, in particular age and full-/part-time status.¹⁵ We refer to area effects estimated this way as 'AKM with controls'.

The third and fourth rows in Table 1 summarise the area effects estimated using these two approaches. The third row (AKM) shows that controlling for time-fixed individual characteristics reduces raw area differences by around a factor of three. Adding controls for time-varying observables slightly reduces estimated differences further, as shown in the fourth row. These 'AKM with controls' estimates suggest that an individual working in the Isle of Wight (the area with the lowest area effects) would, on average, increase their wages by around 17% if they moved to a job in London (with the highest area effects).¹⁶ Individuals moving jobs across most of

¹² Education is not included in the Annual Survey of Hours and Earnings, so we use a measure of skill derived from occupations following Aghion et al. (2019), which roughly corresponds to graduates and non-graduates. See Appendix A for more details.

¹³ After Jacob Mincer, who modelled wages as a function of education.

¹⁴ After Abowd, Kramarz and Margolis (1999), who applied a two-way fixed effects model to decompose the effect of firms and individuals on wages inequality. The assumptions of the AKM model are discussed in Appendix C.

¹⁵ If, for example, people move jobs to work in London when young and when wage growth is high, changes in wages could be attributed to the move rather than to the age. Similarly, if people tend to move out of London and start working part-time when they have children, the change in wages could be attributed to the move rather than to working part-time.

¹⁶ London is the highest-wage area and the area with the highest area effects. However, the area with the lowest 'raw' wages (Thetford and Mildenhall) is not the area with the lowest area effects (Isle of Wight).

the distribution, from an area at the 10th to an area at the 90th percentile, can expect to increase their wages by around 6%.

In short, moving down the rows, spatial differences in average wages look much smaller after controlling for differences between individuals who work in different areas. The compression of the distribution of area wage effects implies that differences in individual characteristics across areas – however they occur – go quite a long way towards balancing the supply and demand of high- and low-skilled workers across areas. From a policy perspective, these results suggest that place-based investments in low-wage areas will do little for average wages in those areas, unless they also change the mix of worker skills, either by upskilling the local population or by attracting skilled workers from elsewhere. We consider the implications for policy in more detail in the final section.

A similar approach can be used to see how spatial differences in employment rates change allowing for differences across areas in individual characteristics. For employment, we do not have panel data so we can only compare raw area differences with the Mincerian area effects that control for observable characteristics. However, because the employment data have much richer information on individual characteristics, we can include more controls than for wages: gender, age, education, ethnicity, whether UK-born, whether a UK citizen and household characteristics interacted with gender (marital status, number of children and age of the youngest child).

The last two rows of Table 1 show the results for employment rates. These show the percentage point difference in employment rates between areas at different parts of the distribution. These are less striking in two ways. First, differences in raw area averages are less pronounced than for wages. For example, the area with the highest employment rate (Basingstoke) has a rate around 16 percentage points higher than the area with the lowest employment rate (Kilmarnock and Irvine). That is a 23% difference in contrast to the 43% difference between the highest- and lowest-wage areas. Second, differences in individual characteristics between areas appear to play less of a role for employment rates than for wages. Controlling for observable differences in individual characteristics reduces employment rate differentials in each of the five measures in Table 1 by around 30% (contrast the bigger reductions seen when moving between the first and second rows for wages).

Differences in area composition matter

Spatial differences in average wages and employment look much smaller allowing for differences in the characteristics of workers across areas. This raises the question of how much of the observed spatial disparities in wages and employment can be attributed to area effects, to the differences in the composition of workers across areas (based on their individual effects) and to the correlation between the two. This subsection provides evidence suggesting that much of the wage variation seen across areas reflects differences in composition rather than differences in the underlying area effects. Area effects play a bigger role for employment, although in part this may be because we can only control for observable characteristics in the employment regressions.

We use the decomposition in Gibbons, Overman and Pelkonen (2014) applied to area averages for wages and employment rates. Detailed results are provided in Appendix C. Box 2 provides a graphical treatment to help with intuition.

After controlling for time-fixed and observed time-variant individual characteristics ('AKM with controls'), just 10% of the variation in average wages across TTWAs is attributed to area effects. 64% of the variation is attributed to differences in individual characteristics – more accurately, variation in area averages of individual effects.

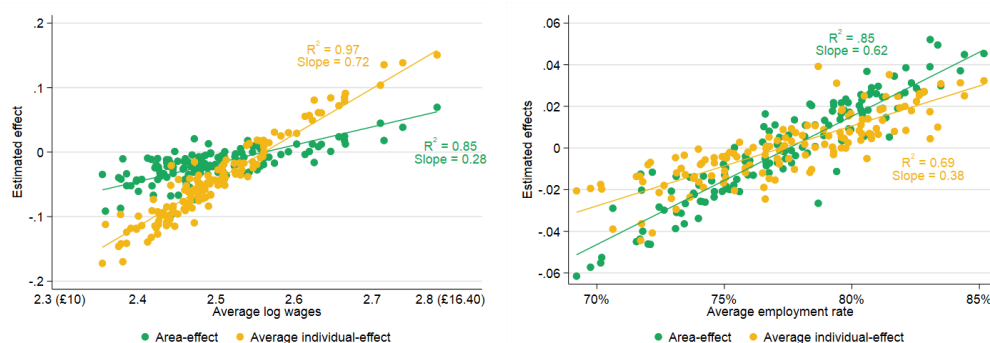
The remaining 26% is attributed to the positive correlation between area and individual characteristics. Whether this is attributed to differences in individual characteristics across areas or to area effects depends on how this correlation is interpreted. One interpretation is that individuals with high earnings potential concentrate in areas that are highly productive (and therefore offer higher wages), in which case the remaining 26% can be thought of as part of the area effect. Another interpretation is that the concentration of high-skilled individuals is what drives high area effects in those areas. If this is the case, the remaining 26% should be attributed to differences in individual characteristics. The reality is likely to lie somewhere between the two extremes.

In short, between 64% and 90% of the wage variation seen across areas reflects differences in composition in terms of the types of workers. Two aspects of this matter. The spatial concentration of workers who would earn higher wages wherever they work plays the biggest role (64%). But it also matters that workers who would earn higher wages wherever they work generally work in higher-paying areas with higher area effects, and vice versa (26%). As just discussed, to the extent that this reflects the concentration of high-skilled individuals in certain areas driving higher area effects, this positive correlation should also be thought of as partly reflecting the role of individual characteristics. Although the available evidence does not allow more precision, it does suggest strong feedback from the concentration of high-skilled individuals to area productivity (Moretti, 2004a; Glaeser and Resseger, 2010), so some of this positive correlation will be due to this concentration.

Box 2. Variance decompositions: intuition

To understand the intuition behind the decomposition, it is helpful to look at two sets of figures. The first, following Card, Rothstein and Yi (2021), plot average wages and employment rates in each TTWA against the estimated area effect and the average individual effect in that area. Panel a of Figure 14 shows estimates for (log) wages controlling for fixed individual characteristics, age and full-/part-time status ('AKM with controls'). Areas with higher average wages have higher area effects, as shown by the green dots. However, average individual effects in the area rise much more steeply with average wages, which implies that much of the variation seen across areas is accounted for by differences in individual characteristics across areas.

Figure 14. The positive correlation between estimated area effects and area averages
(a) Wages (2012–19) (b) Employment rates (2012–19)



Source: Annual Survey of Hours and Earnings (wages); Annual Population Survey (employment).

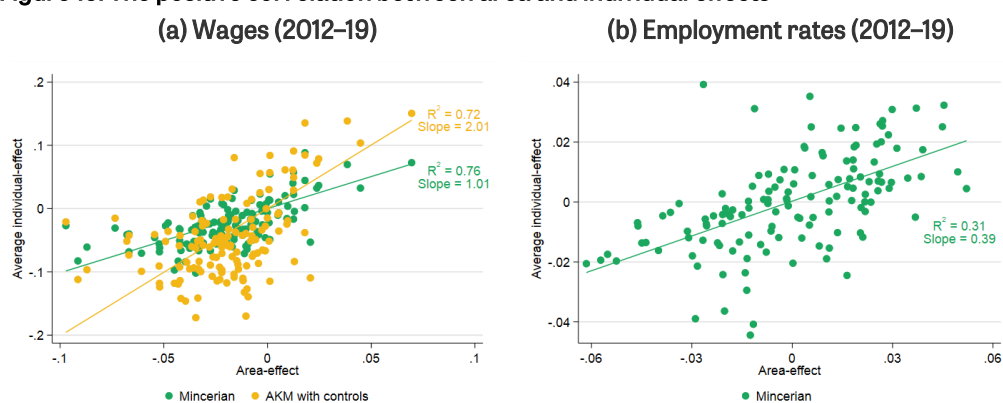
Panel b, which plots results for employment, shows that estimated area effects rise more steeply with employment rates than with wages. This implies that less of the variation seen across areas is accounted for by differences in individual characteristics, and more by the underlying area effects – although, as discussed in the main text, this difference may arise because we can only control for observable characteristics when looking at employment rates.¹⁷

As explained in Card, Rothstein and Yi (2021), the slopes of the lines in these graphs tell us about the relative contribution of area effects and individual effects in accounting for the observed variation in average wages. Consistent with the main text, Figure 14 tells us that for wages, differences in individual-effects across areas play more of a role than area effects, and vice versa for employment. But this comparison does not separate out the correlation between individual and area effects that forms the third part of the decomposition reported in the main text.¹⁸

The fact that all four lines slope up in Figure 14 suggests that areas with high wages have both high area effects and workers who would be highly rewarded wherever they work, and similarly for employment. The second set of figures that help understand the decomposition make this clear by directly plotting estimated area effects against the area-level average of individual effects.

Panel a of Figure 15 shows a positive correlation for wages, meaning that workers with characteristics that are most highly rewarded in the labour market are concentrated in areas where any given worker would command higher wages. The relationship is steeper when time-fixed unobserved characteristics are included in the estimate of individual effects – that is, steeper for the 'AKM with controls' estimates than for the 'Mincerian' estimates. As discussed in more detail in Section 4, this positive relationship is stronger for areas in the top of the distribution of area and individual effects. Panel b shows that this positive correlation also exists for employment, though the correlation is lower than for wages – individuals with a high employment probability (irrespective of where they live) tend to live in areas where any given individual is more likely to find work.

Figure 15. The positive correlation between area and individual effects



Source: Annual Survey of Hours and Earnings (wages); Annual Population Survey (employment).

¹⁷ Consistent with this, Appendix Figure B7 replicates Figure 14(a) for wages using the Mincerian regression and shows that the relative slopes of the two lines are now more comparable to those in Figure 14(b).

¹⁸ Details of the decomposition in Card, Rothstein and Yi (2021) are given in Appendix C.

For employment rates, controlling for observable characteristics (in the Mincerian regression), 46% of the variation across TTWAs is attributed to area effects. Around 20% of the variation is attributed to differences in observable individual characteristics across areas. The remaining 34% is attributed to the positive correlation between area and individual characteristics.

The estimated contribution of differences in individual characteristics across areas to spatial disparities in employment rates is much smaller than that for wages. Partly this difference may arise because we can only control for observable individual characteristics in the employment regression, so some differences in unobserved characteristics are (mis-)attributed to area effects. For comparison, the Mincerian regression for wages attributes 40% of the variation in average wages to area effects, as opposed to 10% when controlling for unobserved time-fixed characteristics ('AKM with controls'). That said, this is still smaller than the 46% estimate for employment and we can control for a much richer set of observable characteristics in the Mincerian employment regressions than in the comparable wage regression. Overall, this suggests that differences in composition likely play a somewhat smaller role in the spatial variation of employment rates – although we cannot be certain how much smaller.

Our results are consistent with results in Knies, Melo and Zhang (2021) in the UK, who apply similar methods on a different data set (Understanding Society) and using different spatial scales ('neighbourhoods' rather than TTWAs). They find that whilst wages are lower in more deprived areas, this entirely reflects differences in the types of people who live in different areas: when individual effects are included, there is no relationship between area deprivation and wages. Further, they show that this result holds using a range of spatial scales – they define neighbourhoods using a range of population thresholds from around 300 people to around 10,000 people.

The results are also in line with findings from other developed countries. Using a method comparable to our 'AKM with controls' and a different decomposition approach, Card, Rothstein and Yi (2021) find that two-thirds of the observed variation in wages across commuting zones in the US is attributable to differences in individual characteristics across areas. Applying the decomposition in Card et al. – see Appendix C for details – to our data attributes just under three-quarters (72%) of the area variation to differences in individual characteristics.¹⁹ The similarity is striking given the size of the US and the amount of variation across US states as compared with the UK. Combes, Duranton and Gobillon (2008), De la Roca and Puga (2017) and Dauth et al. (2018) find that differences in worker characteristics account for a large share of the variation in mean wages and earnings in France, Spain and Germany respectively.

We have focused on the extent to which spatial disparities in labour market outcomes reflect differences in area composition in terms of the characteristics of people who live and work in different places. But the decomposition can also be used to consider the relative roles of individual characteristics and area in understanding individual inequality – something that will be useful when thinking about policy and the consequences of narrowing spatial disparities between areas. Spatial disparities in average wages ('raw') account for 5.3% of the variation in wages across individuals in the post-Great-Recession period. Area effects controlling for observable characteristics ('Mincerian') account for 2.4% of the variation in individual-level wages, and area effects also controlling for fixed characteristics ('AKM with controls') account for 0.5% of the variation. For employment, the corresponding figures are 0.6% for spatial disparities in average

¹⁹ This is larger than the 10% reported above, because the decomposition in Card, Rothstein and Yi (2021) attributes some of the correlation between area effects and individual effects to areas. See Box 2 and Appendix C for further discussion.

employment rates ('raw') and 0.3% when controlling for observables ('Mincerian'). These figures suggest that eliminating area effects would have a very small impact on individual inequalities.

Changes over time: a lot of persistence and differences in composition becoming more important

Section 2 shows that spatial disparities in employment rates and wages are stubbornly persistent. Despite a slight convergence in average wages across areas and somewhat more convergence in employment rates, most areas that were struggling 20 years ago are still struggling today, and vice versa for areas that were flourishing. Appendix Figure C3 shows a similar pattern for average individual effects and for area effects – high persistence for wages, somewhat less persistence for employment. Areas' relative positions – in terms of both the composition of workers in terms of their individual effects and the area effect on wages and employment probabilities – do not move much over time. That said, the graphs also show more movement in terms of area effects on wage and employment probabilities than they do on composition in terms of individual effects. Consistent with this, Figure C4 – which plots the distribution of average individual effects and area effects in the two different periods – shows less convergence in individual effects than in area effects.

This suggests that the extent to which spatial disparities are accounted for by differences in the composition of workers may have increased over time. Once again, we can use variance decompositions to make this statement more precise. As shown in Appendix C, area effects accounted for 13% of the variation in average wages across areas in 1998–2007, and 10% of the variation in 2012–19. This increase in the role of composition is much more pronounced for employment rates, where area effects accounted for 67% of the variation across areas in 1998–2007 and only 46% of the variation in 2012–19.

4. What drives the spatial concentration of high-skilled workers?

Despite modest reductions, overall spatial disparities in labour market outcomes remain large, and areas experience considerable persistence over time. Thinking about area effects and differences in the composition of workers across areas helps understand these patterns. The role of composition is large and possibly growing.

This section considers what drives the spatial concentration of high-skilled workers and the role this plays in the persistence of disparities.²⁰ The distribution of skills across places is an equilibrium outcome reflecting the interaction between the demand for, and supply of, skills. We start by describing the geography of high-skilled jobs and the implications for the demand for skills. We then describe differences in educational attainment and selective migration and the implications for the supply of skills. Next, we show how the demand for, and supply of, skills interact in ways that can be self-reinforcing and hence magnify the effect of shocks and make spatial disparities persistent. We finish by considering the offsetting forces – in particular, differences in the cost of living – that work against these self-reinforcing mechanisms and ensure that not everyone ends up living and working in the most productive areas.

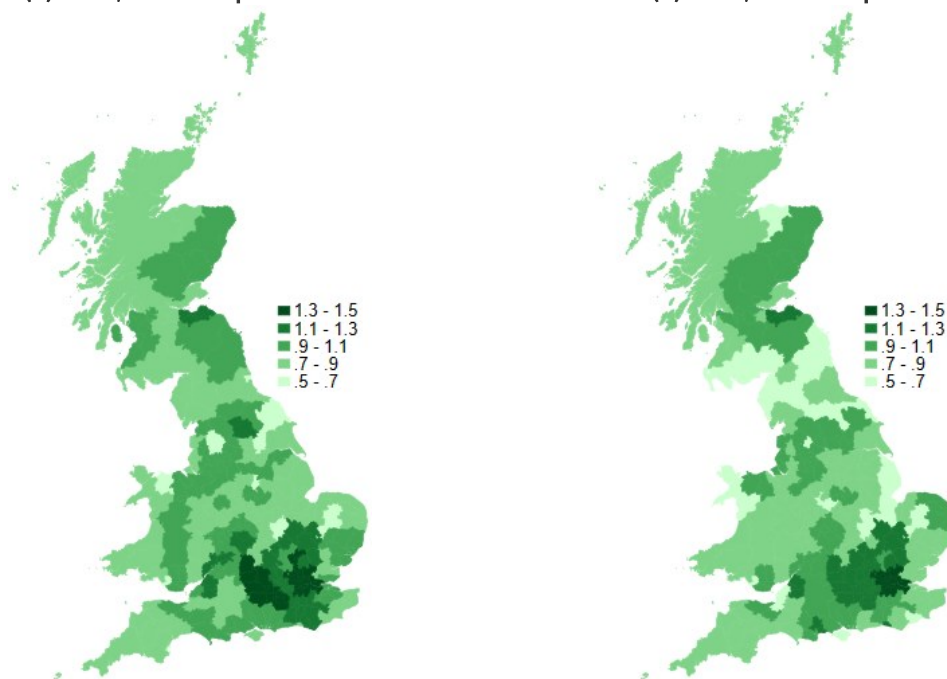
²⁰ We focus on high-skilled workers because, as explained above, it is the spatial concentration of these workers, and of the highest-wage workers in the highest-wage areas, that explains a large share of spatial disparities in labour market outcomes.

The geography of jobs and the demand for skills

One key driver of the spatial concentration of high-skilled workers is the geography of high-skilled jobs. Agglomeration benefits – sharing indivisible infrastructure, the buying and selling of specialist inputs and goods, learning from others – mean that the spatial concentration of firms and workers increases productivity (Duranton and Puga, 2004, 2020). Consensus estimates suggest that doubling population density increases productivity by around 2% (Combes and Gobillon, 2015).²¹ These benefits are larger for high-tech firms and high-skilled workers (Jaffe, Trajtenberg and Henderson, 1993; Glaeser and Resseger, 2010; Combes and Gobillon, 2015), so high-skilled jobs in particular tend to be spatially concentrated to take advantage of these benefits.

Not only are high-skilled jobs spatially concentrated, but they have become more so between 1998 and 2019, as illustrated in Figure 16. To produce these maps, we classify ASHE data on employee jobs using the high and low skill classification from Aghion et al. (2019) used above, where 'high-skilled' jobs roughly correspond to jobs that require a degree. To allow for considerable growth in the national share of graduate jobs, the maps plot the location quotient. A value above 1 means that the area has a higher share of graduate jobs than nationally and vice versa. Compared with 1998, graduate jobs have become less evenly distributed across the country. Appendix Figure B8 shows that measures of the spatial concentration of jobs increased steadily over this period.

Figure 16. High-skilled jobs have become more spatially concentrated between 1998 and 2019
(a) 1998, location quotient (b) 2019, location quotient



Note: Working-age (16–64) population. High-skilled jobs based on classification in Aghion et al. (2019). See Appendix A for more details.

Source: Annual Survey of Hours and Earnings.

²¹ Looking back at Table 1, London's density alone must explain a lot of the 9% difference between the area effect in London and the national average (median).

What explains this increasing spatial concentration? Ongoing globalisation, which has supported offshoring, and technological change that has improved productivity in manufacturing have meant that manufacturing employment has continued to fall (Goos, Manning and Salomons, 2014). The increasing concentration of high-skilled jobs may also reflect a shift towards skill-intensive industries, such as finance and high-tech, which concentrate to take advantage of agglomeration benefits (Glaeser and Resseger, 2010).

Skill-biased technological change has favoured the high-skilled over the low-skilled (Acemoglu, 2002), with differential consequences for areas depending on the skill composition of their workforce.²² Technological changes may also have affected the demand for skills in some areas. Moretti (2013) argues that the dot-com boom in the San Francisco Bay Area was an example of a localised skill-biased shock; similar forces may be at play in parts of the UK.

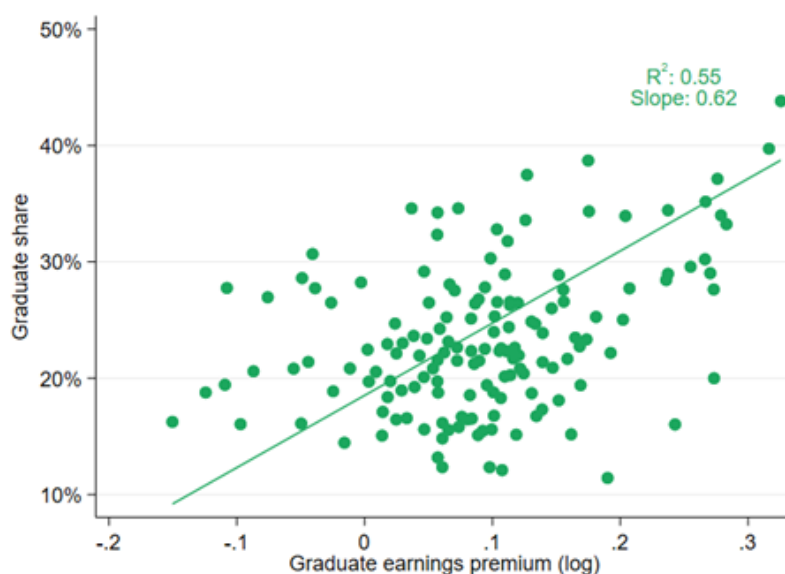
Technological change may also have increased agglomeration benefits for skill-intensive industries (Moretti, 2021). This may explain, for example, the historically high levels of concentration of innovative activity seen today (Kerr and Robert-Nicoud, 2020; Andrews and Whalley, 2021). The benefits from the spatial clustering of skilled workers across all industries may also have increased (Moretti, 2004b).

Increased spatial concentration may also reflect the spatial reorganisation of the economy in response to declining communication and transport costs, which allow firms to serve markets at a distance (Fujita, Krugman and Venables, 1999; Ottaviano and Thisse, 2004; Redding and Turner, 2015). Firms themselves may also have spatially restructured as management practices have changed – in part in response to declining communication and transport costs but also due to technological change. This restructuring often involves the separation of production from management, with high-skilled headquarter jobs located in one city and low-skilled production or back-room activities in another (Duranton and Puga, 2005; Henderson and Ono, 2008; Gokan, Kichko and Thisse, 2019). Finally, the rise of superstar firms (Autor et al., 2020), workers (Rosen, 1981; Alvaredo et al., 2013) and cities (Gyourko et al., 2013) are drivers or consequences of the spatial concentration of high-paying employment and high-paid workers.

These changes explain why the geography of demand for high-skilled workers is spatially concentrated and possibly becoming more so.²³ Moretti (2013) argues that in the US, the concentration of graduates in cities since the 1980s was primarily driven by the relative demand for skills. Figure 17, taken from Britton et al. (2021), provides evidence that the concentration of skills in the UK may also be partly demand-led. It shows a positive correlation across TTWAs between graduate shares in recent GCSE cohorts and the local graduate earnings premium – the percentage difference between graduate and non-graduate earnings in a TTWA, controlling for observable individual characteristics. This suggests that demand for graduates is high in areas with lots of graduates – if this were not the case, competition among graduates would drive down the graduate premium. Appendix Figure B9 replicates the figure using our data, which do not allow for such a precise estimate of local graduate premiums but also show a positive correlation.

²² More recently, economists have focused on the interaction between skills, tasks and technologies in understanding changing inequalities in employment and earnings (Acemoglu and Autor, 2011). While this shift in perspective provides a more nuanced picture of the impact of technology on skills, employment and earnings, ongoing changes still appear to favour areas with a more skilled workforce.

²³ Note that spatial differences in demand for skills are entirely consistent with the low variation in area effects that we find, since the latter is an equilibrium outcome, after individuals have chosen how much education to acquire and where to live based on the opportunities available. Low variation in area effects suggests that differences in area composition have arbitrated away most of the spatial differences in labour demand.

Figure 17. Graduate shares and graduate premiums are positively correlated

Note: Outcomes at age 27 for the 2002–05 GCSE cohorts in England. Local graduate premium estimated using regression of earnings on graduate status interacted with TWA dummies, controlling for differences in detailed measures of prior educational attainment, demographics and socio-economic background.

Source: Britton et al. (2021) using Longitudinal Education Outcomes (LEO) data set.

Educational attainment, selective migration and the supply of skills

While the concentration of skills is partly demand-led, spatial disparities in the supply of skills also play an important role. Differences in the supply of skills reflect two broad factors. First, educational attainment differs for people growing up in different areas. Second, patterns of mobility depend on the level of education and skills. Of course, mobility patterns may differ partly in response to spatial differences in the demand for skills – demand and supply reinforce each other, as discussed in the next subsection.

Figure 18, reproduced from Britton et al. (2021), illustrates spatial differences in both educational attainment and mobility. The figure uses data from the Longitudinal Education Outcomes (LEO) data set for all individuals who completed their GCSEs in England in 2002–05. The horizontal axis shows the share of pupils from an area, at age 16, who have a university degree by age 27. The vertical axis shows the share of pupils who live in an area, at age 27, who have a degree. With no migration, the two measures would coincide, and all points would lie on the 45-degree line.²⁴

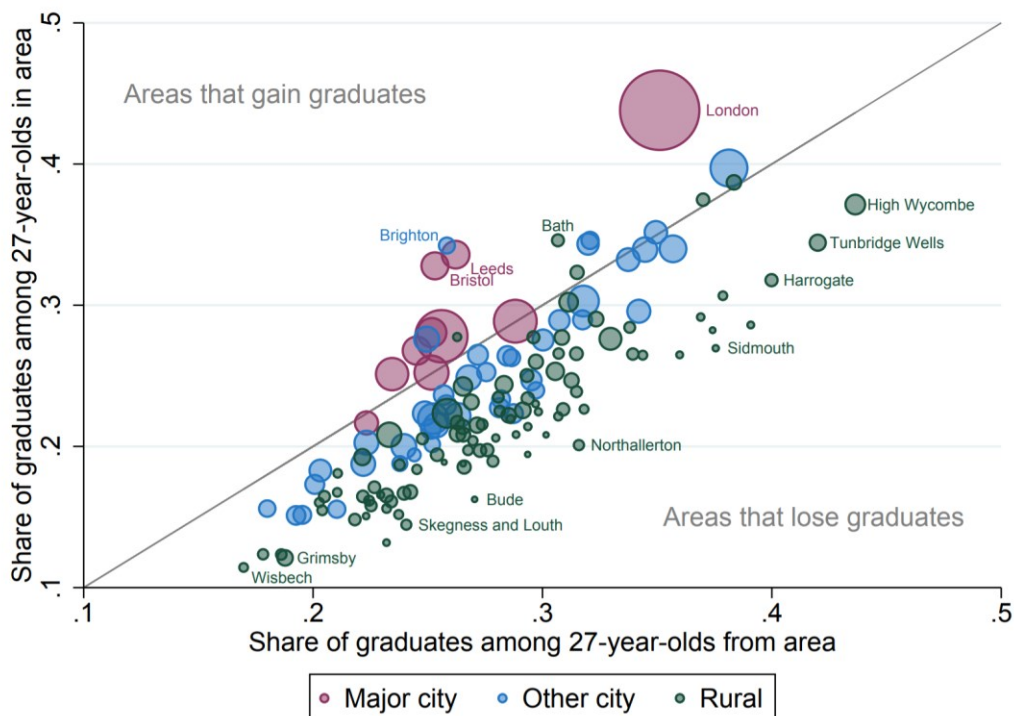
Looking across the horizontal axis, there are substantial differences in educational attainment across areas. Just 19% of children who grew up in Grimsby were graduates by age 27, compared with 42% of children who grew up in Tunbridge Wells. Intergenerational transmission is an important determinant of educational attainment (Black, Devereux and Salvanes, 2005; Fleury and Gilles, 2018), so spatial disparities in educational attainment partly reflect disparities in the

²⁴ A third factor is international migration. Where immigrants of different skill levels choose to live in the UK will affect the local composition of skills. We do not consider international migration in this chapter; this will be explored in a separate chapter by Dustmann, Kostas and Preston (forthcoming).

skills of the adult population. On top of that, there may be area effects on educational outcomes – for example, due to differences in the quality of schools or the composition of local peer groups.²⁵

But this is not the whole story. Graduates from poorer areas tend to leave. By age 27, only 12% of those who live in Grimsby have a degree – half of the 19% of children from Grimsby who got degrees had left.²⁶ In contrast, London, which already has high graduation rates, further attracts graduates through migration. These patterns of selective graduate migration exacerbate spatial disparities in educational attainment. Comparing the concentration of graduates by area of origin at age 16 with the concentration by area of residence at age 27, Britton et al. (2021) find that common measures of concentration increase as a result of migration – the index of dissimilarity increases by around 50% and the Herfindahl–Hirschmann Index (HHI) increases by over two-thirds.

Figure 18. The selective migration of graduates as illustrated by the share of graduates by area of origin versus area of residence, 2013–16



Note: Restricted to 2002–05 GCSE cohorts in England for tax years 2013–14 to 2016–17. Shows outcomes at age 27. Cities correspond to 'Primary Urban Areas' under the Centre for Cities definition. 'Major cities' refers to the 10 largest cities. Areas refer to 2011 TTWAs, not LA-based groups of TTWAs as in the rest of the chapter. The size of the circles represents the age-16 population of each TTWA.

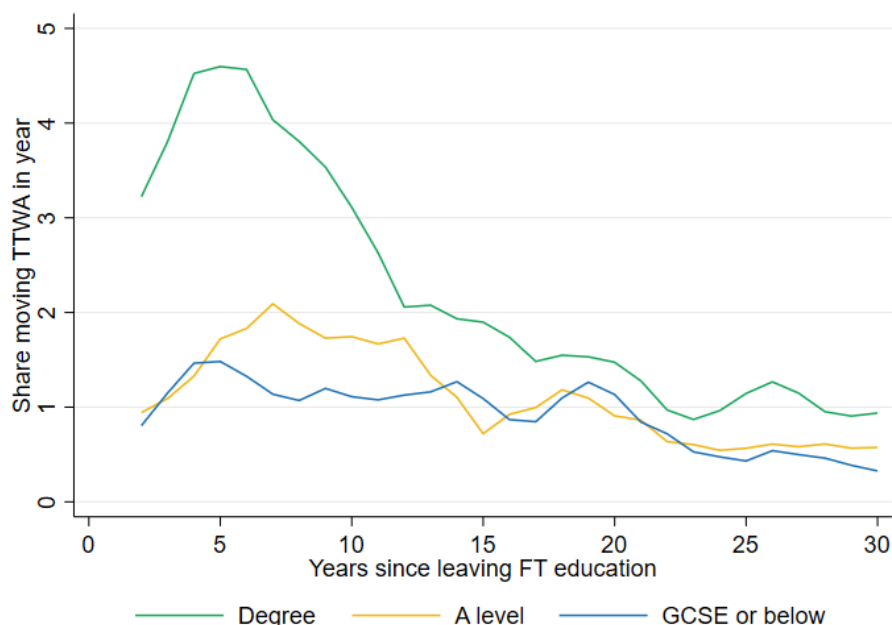
Source: Britton et al. (2021) using Longitudinal Education Outcomes (LEO) data set.

²⁵ We could find no studies that quantify the relative importance of intergenerational transmission and area effects in education in explaining spatial disparities in education across TTWAs or LAs. See footnote 7.

²⁶ This outflow was partially offset by the in-migration of graduates who grew up in other areas.

We described the patterns in Figure 18 as though they reflect selective graduate migration, not migration among non-graduates. This is mainly because – as illustrated in Figure 19 – graduates are much more mobile than non-graduates, especially at the start of their careers.²⁷ Further, conditional on moving, graduates tend to move away from places with poor labour market prospects towards high-skilled and high-wage areas, but this is not true of non-graduates (Britton et al., 2021).

Figure 19. Graduates are much more mobile than non-graduates, especially at the start of their careers



Note: Working-age (16–64) population. Excludes people in full-time (FT) education. Share of individuals who move TTWA in any given year, by education, 2009–19.

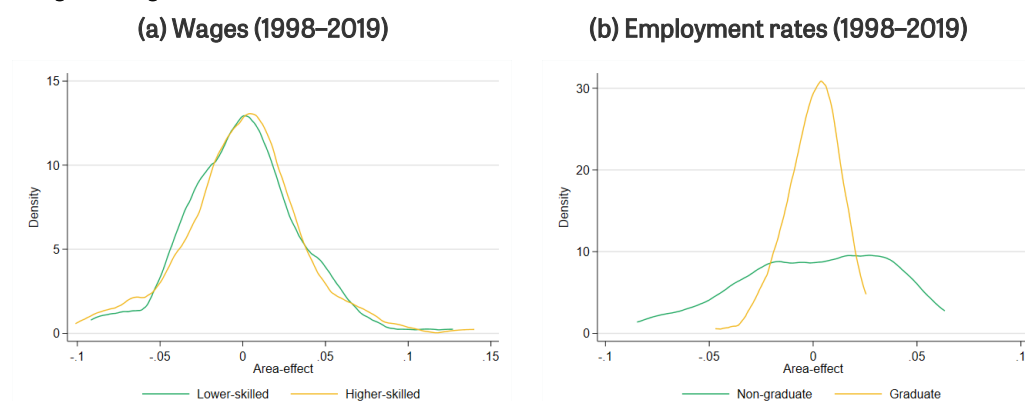
Source: Understanding Society waves 1–10.

Box 3. Why are the high-skilled more mobile?

One possible explanation for higher mobility among the high-skilled is that the gains from moving are larger. Figure 20 suggests that this is not the case for wages and employment rates. Panel a shows that the distribution of estimated area effects for wages looks similar for high- and low-skilled workers, classified using the occupation-based approach from Aghion et al. (2019). This is consistent with Card, Rothstein and Yi (2021), who find in the US that area effects for wages are similar for graduates and non-graduates. Panel b shows that there is more variation in area effects for employment, not less, for non-graduates. Taken together, this suggests that the earnings gains from moving are larger, not smaller, for the low-skilled.

²⁷ Britton et al. (2021) show that Figure 18 is essentially unchanged if they control for non-graduate migration by holding the number of non-graduates in each TTWA constant at age-16 levels.

Figure 20. The distribution of estimated area effects suggests that the earnings gains from moving are larger, not smaller, for the low-skilled



Source: Annual Survey of Hours and Earnings (wages); Labour Force Survey (pre-2004 employment); Annual Population Survey (2004–19 employment).

That said, as shown below, places that offer higher wages also have higher costs of living, including for housing. Because lower-income households spend a larger share of their income on housing (Joseph Rowntree Foundation, 2021), this means that the real gain in earnings (adjusted for costs of living) from moving are likely to be smaller for the low-skilled.

The low-skilled may also face higher psychological, social and financial costs of moving:

- Less-educated workers often relocate without a job in hand because they are employed in sectors and occupations where cross-regional hiring is less common, which makes the move riskier and less appealing. Balgova (2020) estimates that up to half of the education gap in mobility can be attributed to differences between more- and less-educated workers in the likelihood of finding a job in another region.
- Non-graduates are more likely to be in social housing, where rents are subsidised, and it is difficult to find equivalent social housing in a new area. Langella and Manning (2019) find that social renters are less responsive to increases in the local unemployment rate than private renters.
- More speculatively, most graduates in the UK move to a different area to go to university, and this initial move – leaving home for the first time, adapting to a new environment – may make subsequent moves less costly. Britton et al. (2021) find that graduates who moved away for university are, at age 27, more likely to live in a different area which is neither their home town nor the place they went to university.

While we cannot assess the relative importance of these and other factors, the result – as illustrated in Appendix Figure B11 – is that low-skilled individuals are less likely to move for employment reasons, and those who do move primarily do so for family-related or other reasons.

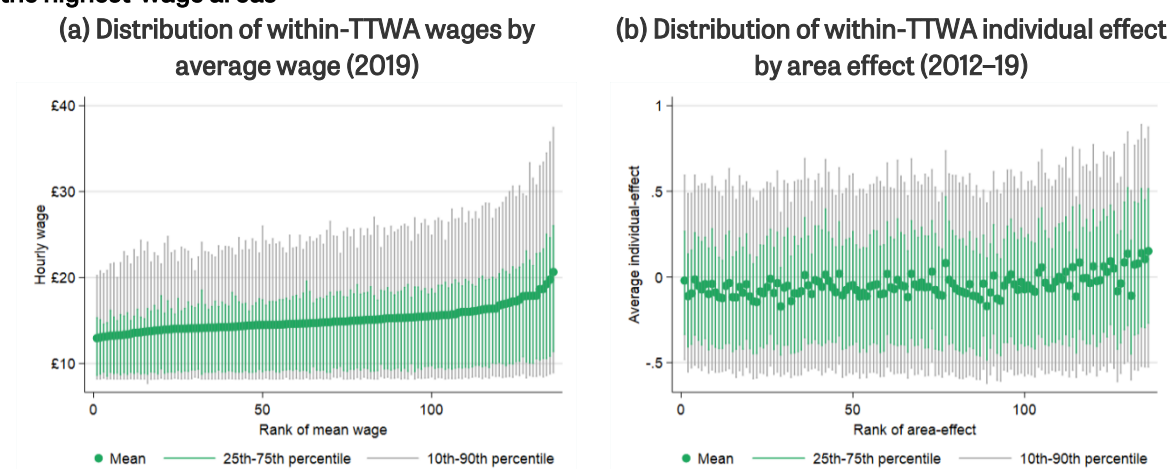
How the demand for, and supply of, skills interact to make spatial disparities large and persistent

Differences in the type of workers across places are an equilibrium outcome that reflects the interaction between the demand for, and supply of, high- and low-skilled workers. We have shown that the geography of high-skilled jobs and the demand for skills are highly spatially uneven. Structural changes, such as the shift towards skill-intensive sectors (which benefit most from agglomeration), the separation of production from management and the rise of superstar firms, mean that the demand for skills is, if anything, becoming more spatially concentrated. The supply of skills also varies across areas, partly due to differences in educational attainment and to selective migration.

These differences in the demand for and supply of skills are self-reinforcing. High-skilled workers from poorer areas may need to move to find suitable employment opportunities. Firms looking to relocate to lower-wage areas struggle to find suitably skilled workers. Firms that need high-skilled workers locate in places with high-skilled workers, and vice versa. This virtuous (or vicious) cycle helps to explain the persistence of spatial disparities.

Panel a of Figure 21, which reproduces Figure 7, suggests that this feedback loop is most pronounced at the top of the wage distribution. As discussed above, it shows that differences in average wages are mostly driven by the top of the wage distribution while there is little variation in wages at the bottom of the wage distribution.

Figure 21. Differences in average wages are mostly driven by the highest-wage people working in the highest-wage areas



Note: Working-age (16–64) population. Area and individual effects in panel b estimated controlling for individual fixed characteristics and time-variant observed characteristics ('AKM with controls'). The vertical axis of panel b is in logs.

Source: Annual Survey of Hours and Earnings.

Panel b of Figure 21 suggests that the positive relationship between area effects and individual effects seen in Figure 15 mainly captures spatial concentration at the top of the wage distribution – workers who would earn the highest wages wherever they worked work in areas with the highest wage premiums as captured by the area effects. The graph ranks TTWAs by estimated area effect and plots the distribution of estimated individual effects within each TTWA. Two things are clear. First, mirroring panel a, we see that differences in the composition of workers across areas are reflected mostly in differences at the top of the wage distribution (note that the vertical axis is in logs, not levels; hence the more compressed distribution than in panel a). Loosely

speaking, 'high-skilled' areas have a larger share of the highest-skilled workers, not a smaller share of the lowest-skilled workers.²⁸ Second, the positive correlation between area and individual effects is more pronounced in the highest-wage areas. Looking across TTWAs, the slope of the green dots becomes steeper for the 10–20 areas with the highest area effects.²⁹

That spatial concentration is most pronounced for the highest-wage people, in the highest-wage areas, helps explain why a small group of areas have wages that are far above the national average. The highest-paid opportunities are concentrated in London and a handful of other areas. To access these jobs, people from less-advantaged areas must 'move out to move on' (Papoutsaki et al., 2020). And this reinforces the tendency of firms that need to access high-skilled workers to locate in this handful of areas. The emergence of amenities that appeal to these highest-wage workers may provide another feedback mechanism that reinforces spatial disparities – discussed in Section 5.

Spatial disparities in the costs of living

To access these spatially concentrated opportunities, workers must pay the costs of living or working in the highest-wage areas. Such costs – alongside congestion and pollution or a preference for natural amenities – are examples of dispersion forces that offset agglomeration benefits. These forces mean that not all households end up located in the highest-wage areas. The same goes for firms economising on rents, taking advantage of lower wages, accessing natural resources or supplying local markets. The balance between agglomeration benefits and dispersion forces partly determines the extent of spatial disparities and how the economy reacts to shocks.

From a household perspective, differences in the cost of living mean that the nominal earnings gains from moving to an area can be offset by higher prices. Figure 22 illustrates this trade-off by plotting estimated area effects in wages against the average (log) private sector rent for two-bedroom properties – using rents as a proxy for the cost of living in the absence of data on local prices.³⁰ Rents vary considerably across TTWAs: the average rent in London, at £1,466 per month in our data, is more than three times the average rent in Burnley. Area effects and rents are positively correlated: places that offer higher wages to a given individual also have higher rents. The estimated slope of 4.8 implies that for a 1% increase in area wage effects, rents increase by 4.8% on average.³¹ Using average rents for three- rather than two-bedroom properties gives a similar slope of 5.0.

²⁸ We do not repeat this exercise for employment probabilities, which are bounded at 1 and therefore not particularly informative. Variation in area-level mean employment probabilities must be driven by differences at the bottom of the distribution.

²⁹ Consistent with this, recall that the positive correlation between area effects and average individual effects accounts for 26% of the variance in mean wages across areas. When we exclude the top 10% of areas when ranked by area effects, this share falls to 15%.

³⁰ The Office for National Statistics occasionally publishes experimental statistics on regional price indices, but these are only available at a very high level of aggregation (NUTS-1 regions), and the last update was for 2016. We use data on average rents at the local authority level, and weight by TTWA-level population to aggregate rents to the TTWA level. We use average rents for two-bedroom properties in our main results; average rents for three-bedroom properties yield very similar results.

³¹ Note that this is the elasticity with respect to estimated area wage effects (using our 'AKM with controls' model), not the elasticity with respect to raw average wages. The latter is 2.25, which is similar to the estimated elasticity in Card, Rothstein and Yi (2021) of 2.70 on quality-adjusted rents.

Figure 22. Differences in the cost of living mean that the nominal earnings gains from moving to an area can be offset by higher prices



Note: Shows average monthly rent for two-bed property in TTWA averaged over 2012–19. Rent data are available at the local authority level, and aggregated to (LA-based) TTWAs based on population size. Area effects estimated controlling for individual fixed characteristics and time-variant observed characteristics ('AKM with controls').

Source: Valuation Office Agency; Welsh Government; Scottish Government; Annual Survey of Hours and Earnings.

A back-of-the-envelope calculation – ignoring social rent, housing benefits, other sources of income and other differences in price levels – implies that if individuals spend more than a fifth of their earnings on housing, the gains from moving to a higher-wage area would be exactly offset by an increase in housing costs. The average share of consumption on housing in 2019–20 in the Living Costs and Food Survey is around 15%.³² Of course, the average rent for a two- or three-bedroom property is only a crude proxy for overall housing costs, and there are other costs-of-living differences besides housing costs. Without a measure of prices at the local level, we cannot compute area effects in wages adjusted for the cost of living. But this back-of-the-envelope calculation suggests that a large share of the gains from moving to higher-wage areas may be eaten up by higher costs of living. This is especially true for low-income households who spend a higher share of their consumption on housing – housing costs make up 17% of total expenditure for households in the bottom fifth of the income distribution, compared with 14% for those in the top fifth. This could be one of the reasons non-graduates are less mobile than graduates, as discussed in Box 3 above.

5. Spatial disparities in well-being

Thus far we have examined spatial disparities in wages, employment and the costs of living. When deciding where to live, people trade off the costs and benefits of different areas (Rosen, 1974; Roback, 1982). Places that offer higher earnings will attract more people, which will in turn push up housing costs, explaining the positive correlation between earnings and housing costs. As well

³² This includes rent net of benefits, mortgage interest payments and council tax.

as real incomes, however, people also care about the amenities that different places have to offer – congestion, pollution, the natural environment, the presence of restaurants and shops and so on. If people are sufficiently mobile, the utility an individual derives from living in an area – taking all these factors into account – should be broadly equalised across areas.

Although individual utility is not directly observable, we can look at differences in self-reported well-being across areas using two common measures – life satisfaction and happiness. Summary statistics for these two measures are shown in Table 2. The table reports raw differences between areas as well as differences controlling for observable demographic characteristics known to be correlated with well-being, analogous to the regressions used for wages and employment.

Table 2. There is little variation in self-reported well-being across areas

	Mean	Standard deviation	Max–Min	Med–Min	Max–Med	P90–P10	P75–P25
Life satisfaction (2013–19)							
Raw	7.54	1.77	0.55	0.23	0.32	0.24	0.14
Mincerian			0.46	0.16	0.31	0.19	0.08
Happiness (2013–19)							
Raw	7.80	1.69	0.54	0.21	0.33	0.27	0.13
Mincerian			0.46	0.16	0.30	0.20	0.09

Note: Working-age (16–64) population. Each measure is reported on an 11-point scale ranging from 0 to 10. Well-being questions were first asked in the APS in 2013, and data are pooled from 2013 to 2019.

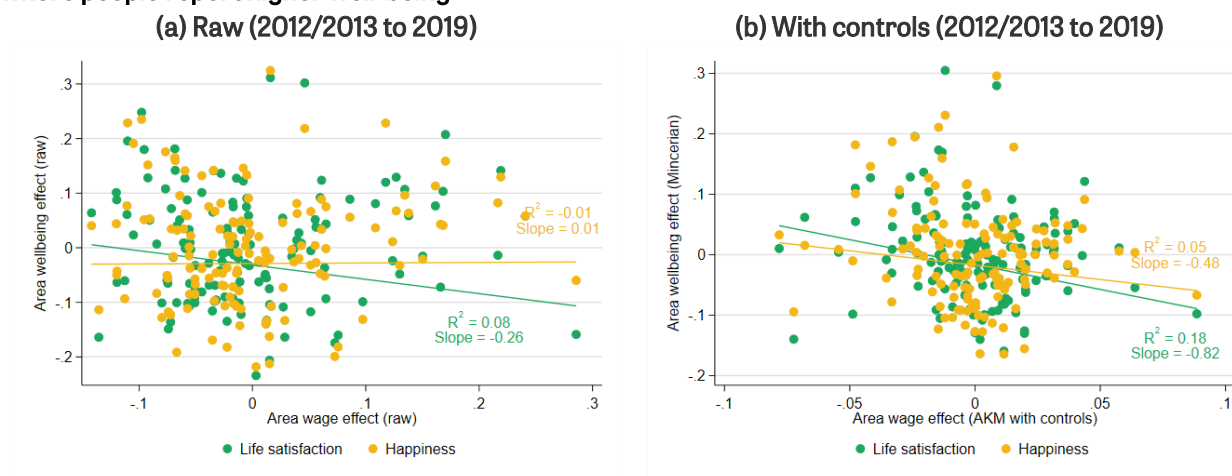
Source: Annual Population Survey.

For life satisfaction, the raw difference between the highest- and lowest-ranked TTWAs (Inverness and Liverpool respectively) is 0.55 points on an 11-point scale or a third of a standard deviation. Controlling for demographic characteristics reduces this to 0.46 points. The difference across most of the distribution – between the 90th and 10th percentiles – is 0.24 points, falling to 0.19 points when controlling for demographic characteristics. The variation in self-reported happiness across areas is similar. For comparison, the difference in average self-reported life satisfaction scores between those with degrees and those without is 0.24 points.

In addition to comparing with the effect of individual characteristics, such as having a degree, another way to interpret the degree of variation in well-being across areas is to compare it with variation in wages. Area disparities in mean wages accounted for 5.3% of the variation in individual wages in the post-Great-Recession period, and area effects estimated using a Mincerian regression accounted for 2.4% of the variation in individual wages. In contrast, area disparities in life satisfaction accounted for just 0.5% of the variation in life satisfaction across individuals, and 0.2% controlling for demographic characteristics. The corresponding figures for happiness are 0.2% and 0.1% respectively. These comparisons tell us that differences in well-being between areas are small – at least compared with differences in wages – and that spatial disparities explain a negligible part of the variation in well-being across individuals.

As well as showing that spatial disparities in well-being are small compared with disparities in labour market outcomes, two other findings are of interest from a policy perspective. First, Figure 23 shows that the correlation between average wages in an area and average self-reported well-being is not positive and may be slightly negative – especially controlling for individual characteristics. Second, Appendix Figure B12 shows that the lack of a positive correlation between wages and well-being holds for both graduates and non-graduates, even though non-graduates are much less mobile than graduates.

Figure 23. Places where people have better labour market outcomes are not generally places where people report higher well-being



Note: Working-age (16–64) population. Uses data on wages from 2012 to 2019 and data on well-being from 2013 to 2019. Regression line weighted by TTWA-level population size.

Source: Annual Population Survey; Annual Survey of Hours and Earnings.

Our results are consistent with Knies, Melo and Zhang (2021), who use panel data to examine the effect of neighbourhood on life satisfaction. They find that once unobserved time-invariant characteristics are controlled for, there is no correlation between neighbourhood deprivation and life satisfaction, using a range of neighbourhood scales from around 300 people to 10,000 people per neighbourhood. Knies and Melo (2019) also find no correlation between neighbourhood deprivation and physical or mental health-related quality-of-life measures once controlling for individual effects.³³ In short, while inequality in well-being is important and there are many reasons to want to improve labour market performance in struggling areas, it should not be assumed that policies targeted in this way will also help target places with low self-reported well-being.

Amenities and the location choices of workers

The canonical model of spatial economics predicts that in equilibrium, utility is equalised across areas, and the cost of living in an area reflects local wages and amenities (Rosen, 1974; Roback, 1982). Building on this insight, urban economists use the difference between area wage effects and costs of living to infer local amenities (Albouy, 2016). This approach gives an idea of the quality of life that places offer while avoiding the need to collect data on lots of different amenities

³³ These measures are derived from the 12-item Short Form Health Survey (SF-12), a self-reported assessment of health relating to eight dimensions of physical functioning, role limitations due to physical and emotional health problems, freedom from bodily pain, general health perception, vitality, social functioning and mental health. These self-assessments are weighted to calculate a physical and mental component summary of health-related quality of life.

and apply a set of arbitrary weights to those different measures to construct a quality of life index.

The basic idea is simple. In places with high housing costs relative to wages, something – which we call high amenities – must make people willing to pay rents which are high relative to wages. In areas with high wages relative to housing costs, something – which we call low amenities – is compensated by the fact that rents are cheap relative to wages. Of course, if the economy is far from equilibrium – for example, because workers do not move to equalise utility – then the index constructed this way will be misleading. As shown below, however, this approach gives a measure of quality of life across areas that fits with popular understanding as well as the demographic differences that were documented in Section 2.

As discussed earlier, we do not have a measure of price levels at the local level, but average rents can be used as a rough proxy for local prices. We therefore use the relationship between wages and rents – plotted in Figure 22 – to derive a proxy for local amenities, measured by the difference between actual and predicted rents as represented by the upward-sloping line in the figure. Areas that have rents that are relatively high given the extra wages that someone earns by living there – that is, areas above the fitted line – are interpreted as having high amenities and vice versa.

Table 3. Areas with lowest and highest implied amenities, 2012–19

Lowest		Highest	
1	Barnsley	1	Brighton
2	Blyth and Ashington	2	Bath
3	Durham and Bishop Auckland	3	London
4	Burnley	4	Eastbourne
5	Rhyl	5	Slough and Heathrow
6	Grimsby	6	Southend
7	Hull	7	Torquay and Paignton
8	Ayr	8	Chichester and Bognor Regis
9	Scunthorpe	9	Crawley
10	Elgin	10	Isle of Wight

Note: Implied amenities defined as difference between actual rents and predicted rents based on unweighted regression of average rents for two-bedroom property on AKM area wage effects in 2012–19. Area effects estimated controlling for individual fixed characteristics and time-variant observed characteristics ('AKM with controls').

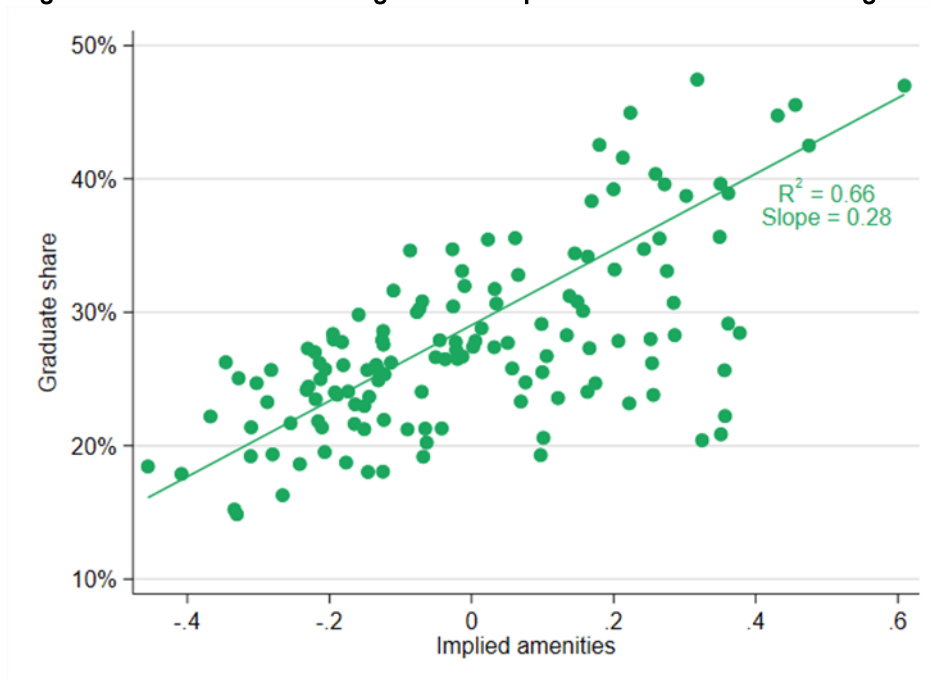
Source: Valuation Office Agency; Welsh Government; Scottish Government; Annual Survey of Hours and Earnings.

Table 3 lists the 10 TTWAs with the lowest and highest amenity levels estimated using this method. Areas with high implied amenities include cities as well as seaside locations. Many ex-industrial towns fall in the list of areas with the lowest levels of implied amenities.

Figure 24 shows that areas that have a high level of implied amenities also have a high share of graduates. This suggests that the concentration of skilled workers in certain areas may partly be driven by preferences for amenities, as well as by labour demand and supply as discussed in Section 4. There may also be important interactions. Graduates may choose to live in an area for

its amenities and this may attract firms that want to employ them. Alternatively, graduates may locate in areas in response to demand for skilled labour, but amenities such as restaurants and theatres are then set up to cater to those graduate populations, making these areas more attractive to high-skilled people. The interaction between labour demand and amenities provides another feedback mechanism which can make spatial concentration self-reinforcing.

Figure 24. Areas that have a high level of implied amenities also have a high share of graduates



Note: Working-age (16–64) population, excluding those in full-time education. Implied amenities defined as difference between actual rents and predicted rents based on unweighted regression of average rents for two-bedroom property on AKM area wage effects in 2012–19. Regression line weighted by TTWA-level population size.

Source: Valuation Office Agency; Welsh Government; Scottish Government; Annual Survey of Hours and Earnings; Annual Population Survey.

The level of local amenities, and the trade-offs places offer in terms of wages, amenities and costs of living, also help explain the demographic patterns seen in Section 2. London and Brighton – cities with high amenities and high wages – both have a high share of graduates and relatively young populations. Older people are concentrated in coastal and rural areas such as Eastbourne, Torquay and Paignton, and Chichester and Bognor Regis, which offer high amenities but relatively low wages.

6. Conclusions and policy implications

Many things determine spatial disparities in Britain. The legacy of 1970s deindustrialisation, the ongoing shift from manufacturing to services, falling communication and transportation costs and the spatial reorganisation of the economy, in response to these and other changes, all play a part in changing the geography of jobs and the demand for different types of workers. Spatial differences in educational attainment, the selective migration of skilled workers and differences in amenities and costs of living help determine the supply of different types of workers. Demand for and supply of skills interact in a way that can be self-reinforcing, meaning large spatial differences can emerge and persist.

One important consequence is differences in the distribution of high- and low-skilled workers across areas. Workers who would earn higher wages wherever they work are spatially concentrated and tend to work in the highest-paying areas (and vice versa). The spatial concentration of skills takes individual labour market advantages and magnifies them further.

All of this has important consequences for many aspects of policy. In this conclusion, we outline some key messages for policies aiming to narrow economic disparities between areas – 'levelling up' in the language of the current government. We are not exhaustive but rather aim to outline some of the ways in which our analysis can be used to consider the impact of different policy options.

Our analysis makes it clear that successfully delivering 'levelling up' will require further clarity on whether levelling up applies to people or places, and on which outcomes are to be levelled up. Is the goal to improve outcomes for people who currently live in left-behind areas or to reduce disparities between places? And what is to be levelled up – employment, education, well-being, 'pride in place' or something else? It will be difficult to develop policy without answering these questions, because the policies that are most likely to be effective are different depending on the answers. For example, if the aim of levelling-up policy is to improve outcomes for people in left-behind places, one policy solution might be to enable more graduates from these areas to move to thriving cities such as London. But such a policy would not serve to equalise outcomes between places. Conversely, targeted investments in certain areas might help to equalise area-level outcomes by attracting high-skilled workers from elsewhere, but this will not necessarily help low-income people living in the area, as discussed more below.

The evidence suggests two important constraints on any policy aiming to reduce disparities between labour markets. The first comes from the fact that over 40% of people only ever work while living in the local labour market where they were born (Bosquet and Overman, 2019). This suggests that policy needs to be realistic about how far people will move for work, particularly for less-educated workers (the figure rises to 50% for those without a degree). Having 'everyone' move to access good jobs is not economically feasible, nor socially or politically acceptable.

The second constraint is true for the other extreme: achieving a level playing field where productivity is equalised and high-skilled jobs are evenly spread. Realism is needed about the market forces at work. Equal outcomes across places would require places to have similar skill compositions and to be of similar sizes. As with the previous strategy, this is not economically feasible, nor socially or politically acceptable. It would also come at a substantial cost in terms of national economic performance, as it would require undermining many of the productivity benefits generated by the concentration of high-skilled workers in London and the South East.

A pragmatic aim might be to improve economic performance in some areas outside of London and the South East – reducing spatial disparities at the regional level, if not necessarily across more narrowly defined local labour markets. This would allow talented young people in left-behind places to access opportunities without having to move across the country. To generate these opportunities and counter the self-reinforcing feedback loops we have described, large investments will be needed in a limited number of places to attract high-skilled workers as well as productive firms. Given that these policies are likely to benefit high-skilled workers more than low-skilled workers, complementary investments will be needed to improve educational attainment and tackle other barriers to working. The rest of this section considers these options in more detail.

Place-based policies

There are many interventions discussed as options for 'levelling up' and generating more economic opportunities outside London and the South East. For example, on some measures, London receives a disproportionate share of infrastructure investment. Many argue that this is a driver of spatial disparities in economic performance, and that investment in infrastructure – particularly in transport – should be more evenly spread across the country. Thinking through the impacts of transport investments provides a good example of where our analysis sheds light on the likely economic impacts.

Infrastructure matters for economic performance. London, and other big cities, could not function without their underlying infrastructure. Commuter towns and rural areas rely on transport infrastructure to connect residents to jobs elsewhere. Firms in all areas rely on transport infrastructure to serve their markets. Evidence suggests that the UK has under-invested in transport infrastructure and that increased investment would improve national economic performance (Aghion et al., 2013).

However, our analysis suggests that the impact of transport investments on labour market disparities will be small unless they significantly alter the composition of the workforce in an area. This is because, as seen in Section 3, differences in area effects (which will capture differences in infrastructure) play a relatively small role in explaining disparities, compared with differences in education and skills. To have a large impact on spatial disparities, transport investments need to make places attractive to high-skilled workers or to the firms that employ them, and to see their skill composition shift accordingly.

At smaller spatial scales, such effects are often seen – for example, if a new railway station makes a neighbourhood attractive for households looking for cheaper housing and easier commutes. At larger spatial scales, however, for infrastructure to have a big effect on spatial disparities in economic performance it must lead to the relocation of large numbers of skilled workers and the firms that employ them, away from London and the South East. Even a project of the size of HS2, for example, will do little for the economy of the West Midlands unless it somehow improves local educational outcomes for children growing up there or encourages a much larger share of graduates and the firms that employ them to locate there.

Infrastructure is not the only example of an intervention whose effects on spatial disparities will depend mainly on the extent to which it changes the skills composition of an area. For example, targeted investment in research and development might be used to help support local innovation rates. Or public sector relocation might create new jobs in an area. Policy can also directly increase local employment opportunities by paying private sector firms to do so (Criscuolo et al., 2019). In all these cases, unless the targeted area is quite small, the direct employment impacts are likely to be small relative to the overall size of the local labour market. Larger local economic benefits will need the intervention to make places attractive to high-skilled workers or to the firms that employ them, and for the area to see its skill composition shift accordingly.

Agglomeration matters

Place-based policies of the kind discussed above could improve labour market outcomes in small towns. But there are many small towns, investment in infrastructure and innovation is costly, and there is a limit to the number of public sector jobs or private sector firms that can be relocated. The approach does not scale up to produce large effects across lots of areas.

Despite these limitations, the politics of levelling up will push for spending to be spread across areas. From an economic perspective, this creates a challenge, because counteracting the economic forces polarising Britain – in particular, the concentration of high-skilled workers and the firms that employ them in our largest city – is likely to need large, spatially focused investments. Such concentrated investments are the most realistic way to generate significant productivity improvements for large numbers of firms and improved employment opportunities for large numbers of people.

As discussed in Section 4, the advantages of high-skilled areas are self-reinforcing. The concentration of high-skilled firms and workers generates productivity advantages for firms and better labour market outcomes for workers. In turn, these better labour market outcomes attract high-skilled workers from across the country. And as seen in Section 5, graduates are attracted to London by more than just economic opportunities. High house prices relative to incomes point to substantial 'amenity benefits' that make London attractive to those who can afford it. In short, London's economic advantages stem from the concentration of skilled firms and workers, and from its economic size, and these factors are self-reinforcing.

London's economic strength also spills over to benefit other towns and cities across the wider South East. Many of London's high-skilled workers commute into the city – to benefit from the employment opportunities it offers – while taking advantage of the different amenity and cost-of-living trade-off offered by commuter towns. The resulting concentrations of high-skilled workers can also attract firms that want to employ such workers but without paying the high rental costs associated with central London offices.

To provide a counterbalance to London and the South East, investment needs to kick-start these self-reinforcing processes elsewhere. Given that size is one key part of this self-reinforcing cycle, this will require major investment – for example, in infrastructure, in R&D and in housing – that is spatially focused on a limited number of places of sufficient size.

Mobility and housing

Policy may also want to support increased mobility to help people access the opportunities created by investment. For example, could policy widen the horizons of young people growing up in disadvantaged areas, so they are willing to commute or move to access opportunities offered in the wider region? Funding mechanisms might be considered that would widen the horizon of young adults going through the further education system in the same way as the university system appears to encourage the mobility of graduates (Britton et al., 2021).

Solutions of this 'on your bike' kind are often highly controversial for reasons that parallel some of the concerns about selective graduate mobility – why should people be 'forced' to move? One answer, as we have seen, is that strong market forces – particularly for high-skilled jobs – may make it difficult for policy to bring jobs to people. Setting this issue aside, while increased mobility would tend to lower spatial disparities in average wages, would it do much good for the low-skilled workers involved? Although the evidence suggests that the wage gains from such mobility are quite small, the employment gains are larger. But Appendix Figure B10 suggests that area wage effects for low-skilled workers are highly correlated with those for high-skilled workers, which means that accessing them requires moving to areas where the high-skilled are concentrated. For lower-paid workers, who spend a high share of their income on housing, the moderate earnings benefits from such moves could be entirely swamped by higher housing costs, as shown in Figure 22.

None of this is to say that policy should not think about ways to support mobility of low-skilled workers. But it needs to be realised that housing costs reduce, or even eliminate, the gains they will experience when moving to more productive areas.

This highlights another policy lever: increasing the supply of housing in more productive areas. As has been extensively documented, more responsive housing supply would help address Britain's chronic affordability problem (Hilber and Vermeulen, 2016; Cheshire and Hilber, 2019). Unfortunately, repeated failures to tackle the UK's housing supply problems suggest that increasing the supply of housing in more productive area is easier said than done.

Helping those at the bottom of the income distribution

This is doubly unfortunate because, while Britain's unaffordable housing may help partly explain the selective migration that underpins spatial disparities, it also has disastrous implications for those at the bottom of the income distribution. As Figure 7 makes clear, differences in wages between areas are small at the bottom of the distribution. Wages at the 10th percentile are very similar across areas – not surprising, since the minimum wage is the same everywhere – and differences at the 25th percentile are also small. The combination of low wages at the bottom and high housing costs means that, for poorer households in particular, incomes adjusted for living costs are lower in higher-wage areas. This is one reason why London has high poverty rates compared with the rest of the country (Agrawal and Phillips, 2020).

All of this suggests that creating counterbalances to London and the South East that decrease spatial disparities by reducing the spatial concentration of high-skilled jobs and workers will tend to benefit higher-wage rather than lower-wage workers. Some of these benefits will trickle down to the lower-paid in the form of moderately higher wages and improved employment rates, but at the cost of more expensive housing. Other non-pecuniary benefits are possible. For talented children growing up in struggling towns, increased opportunities nearby offer the option of commuting or a small-distance move, making it easier to maintain links with family and friends left behind. Improved peer effects may improve educational outcomes and a higher graduate premium may improve incentives to invest in education.

Sadly, while all these trickle-down benefits are possible, London – with its many poor neighbourhoods – points to the limits of this approach for improving outcomes for those at the bottom of the income distribution. This observation has profound implications for rebalancing strategies that focus on reducing selective mobility and increasing the range of options for graduates. Just as seen in London, a 'Northern Powerhouse' or 'Midlands Engine' may do little to improve living standards for low-skilled workers 'fortunate' enough to live there. A more equal spread of graduates will help reduce spatial disparities and may even help improve the overall performance of the economy, but it is no simple fix for improving outcomes for low-skilled workers. To do this, complementary investments will be needed to make sure that households can access the opportunities generated.

The current debate often interprets this as being about 'better transport'. For many poorer households, however, transport investment generally will not be enough. Again, examples from London illustrate the limits to this approach – Barking and Dagenham (areas in the East of London) have good transport links to one of the largest concentrations of employment in the world, but this is not enough to prevent poor outcomes for many households who live there. If poorer households are to benefit from the kind of investments described above, then they will need help to improve their education and skills.

For some households, the multiple barriers that prevent individuals from being able to access better economic opportunities go beyond education and skills. Many of the 'left-behind' places that levelling up may want to target have high proportions of vulnerable people with complex needs and low levels of economic activity. This compounds their problems, as long-term unemployment, poverty, mental illness and poor health often go hand-in-hand.

Addressing these multiple barriers will involve significant investment not only in education and skills, but also in childcare, and in mental and physical health services. Research suggests that small tinkering and minor tweaks of existing policies will not be enough to tackle the multiple barriers faced in these places.

We have focused on the economics of levelling up but it is important to be clear that spending on levelling up does not always need to be justified based on economic growth. There are important public good arguments that can justify increased expenditure across a wide range of policy areas. For example, it is possible to argue for subsidising rural broadband as a public good while recognising that its economic impacts are likely to be limited. Although these policies may not be specifically targeted at the bottom of the income distribution, they will often benefit poorer families most.

An alternative to a narrow economic focus would be to consider well-being. While improving well-being might provide an alternative justification for levelling up, as we have shown, there is no positive correlation between self-reported well-being and economic performance. Indeed, a well-being case for levelling up would suggest considerable funds need to be spent in London – which hardly fits with the current political narrative. This is not the only example where the evidence suggests a pattern of spending that may be a difficult political sell. For example, while the economic case for the need for spatially concentrated investments to level up is strong, the political case is weaker (especially for constituency-based politicians).

The place for place-based policies

Places matter to people. For many people, the place where they grow up will become the place where they live and work. The disparities that we have documented – in labour market opportunities, in costs of living and in amenities – provide the context for, and directly influence, the decisions they take and the life they will live. More formally, spatial disparities are important because local social and economic conditions affect individual outcomes.

But it is possible to overstate the importance of place for determining social and economic outcomes and life chances more generally. Spatial disparities also reflect individual inequality. Growing up in a poor place is more likely to mean growing up in a poor family, and the evidence suggests that the effect of place is small compared with the effect of family. Again, more formally, far more of the variation in individual outcomes is explained by individual characteristics than by area.

The link between individual and spatial disparities is complicated by the fact that, while many people stay close to where they grow up, many others move around. This matters for thinking about what spatial disparities can tell us about important policy issues. For example, the geography of the Brexit vote was highly uneven, with some places more likely to vote Leave and others more likely to vote Remain. One explanation is that the Leave vote reflects the 'revenge' of 'left-behind' places – that is, it is a story not about individuals, but about shared anger by those living in places left behind by technological change and globalisation. The alternative is to think of this as a story about individuals, left behind by the same forces, and where they live.

The first way of thinking about this appears to be driving the current policy response. But the second is perhaps a more useful way of understanding why wealthy Sevenoaks and struggling Sunderland both voted Leave. Different kinds of people, with very different concerns about the European Union, living in different places – but agreeing on the same solution.

Individual mobility also matters because it means that policies that are place-based – that is, targeted at specific places – do not necessarily end up benefiting the people that they aim to help. As we have argued, for example, transport improvements in a poorer area do not necessarily end up benefiting poorer families if improvements in labour market outcomes are small but increases in house prices and rents are large. This is a specific example of a more general problem – place-based policies are a blunt tool. For many policies, using an area-based approach to targeting policy will fail to address within-area disparities or reduce individual inequalities that, as we have shown, are much bigger than the between-area disparities that these policies seek to address. And sometimes policies that could address individual inequality – including spatially focused investments to generate opportunities and spending to help disadvantaged families access them, as discussed above – may have the opposite effect on spatial disparities.

Spatial disparities in the UK are profound and persistent. Improving economic performance and helping to tackle the problems of left-behind places are both important policy objectives. Addressing these challenges requires a new approach to policy, one that allows for different responses in different places.

Such variation makes many people nervous. But we would argue that we should care more about the effect of policies on people than on places. If this is the case, policies should be judged on the extent to which they improve individual opportunities and on who benefits, rather than on whether they narrow the gap between places.

References

- Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999), 'High Wage Workers and High Wage Firms', *Econometrica*, 67(2), 251–333, <https://www.jstor.org/stable/2999586>.
- Acemoglu, D. (2002), 'Technical Change, Inequality, and the Labor Market', *Journal of Economic Literature*, 40(1), 7–72, <https://doi.org/10.1257/0022051026976>.
- Acemoglu, D., and Autor, D. (2011), 'Skills, Tasks and Technologies: Implications for Employment and Earnings', in O. Ashenfelter and D. Card (eds), *Handbook of Labor Economics*, 4B, Elsevier.
- Aghion, P., Bergeaud, A., Blundell, R., and Griffith, R. (2019), 'The Innovation Premium to Soft Skills in Low-Skilled Occupations', Centre for Economic Performance, Discussion Paper 1665, <https://cep.lse.ac.uk/pubs/download/dp1665.pdf>.
- Aghion, P., Besley, T., Browne, J., Caselli, F., Lambert, R., Lomax, R., Pissarides, C., Stern, N., and Van Reenen, J. (2013), 'Investing for Prosperity: Skills, Infrastructure and Innovation', Centre for Economic Performance, Special Paper.
- Agrawal, S., and Phillips, D. (2020), 'Catching Up or Falling Behind? Geographical Inequalities in the UK and How They Have Changed in Recent Years', Briefing Note, IFS Deaton Review of Inequalities, <https://ifs.org.uk/publications/14969>.
- Albouy, D. (2016), 'What Are Cities Worth? Land Rents, Local Productivity, and the Total Value of Amenities', *Review of Economics and Statistics*, 98(3), 477–87, https://doi.org/10.1162/REST_a_00550.
- Alvaredo, F., Atkinson, A. B., Piketty, T., and Saez, E. (2013), 'The Top 1 Percent in International and Historical Perspective', *Journal of Economic Perspectives*, 27(3), 3–20, <https://doi.org/10.1257/jep.27.3.3>.
- Andrews, J. (2017), 'Analysis: The Introduction of Progress 8', Education Policy Institute, <https://epi.org.uk/publications-and-research/analysis-introduction-progress-8/>.
- Andrews, M., and Whalley, A. (2021), '150 Years of the Geography of Innovation', *Regional Science and Urban Economics*, forthcoming, <https://doi.org/10.1016/j.regsciurbeco.2020.103627>.
- Autor, D., Dorn, D., Katz, L., Patterson, C., and Van Reenen, J. (2020), 'The Fall of the Labor Share and the Rise of Superstar Firms', *Quarterly Journal of Economics*, 135(2), 645–709, <https://doi.org/10.1093/qje/qjaa004>.
- Balgova, M. (2020), 'Leaping into the Unknown? The Role of Job Search in Migration Decisions', <https://mariabalgova.github.io/JMPBalgova.pdf>.
- Benson, R., Duffy, B., Hesketh, R., and Hewlett, K. (2021), 'Attitudes to Inequalities', IFS Deaton Review of Inequalities, <https://ifs.org.uk/inequality/wp-content/uploads/2021/09/Attitudes-to-inequalities.pdf>.
- Black, S. E., Devereux, P. J., and Salvanes, K. G. (2005), 'Why the Apple Doesn't Fall Far: Understanding Intergenerational Transmission of Human Capital', *American Economic Review*, 95(1), 437–49, <https://doi.org/10.1257/0002828053828635>.

- Blundell, R., Duncan, A., McCrae, J., and Meghir, C. (2000), 'The Labour Market Impact of the Working Families' Tax Credit', *Fiscal Studies*, 21(1), 75–104, <https://doi.org/10.1111/j.1475-5890.2000.tb00581.x>.
- Blundell, R., Joyce, R., Norris Keiller, A., and Ziliak, J. (2018), 'Income Inequality and the Labour Market in Britain and the US', *Journal of Public Economics*, 162, 48–62, <https://doi.org/10.1016/j.jpubeco.2018.04.001>.
- Bosquet, C., and Overman, H. (2019), 'Why Does Birthplace Matter So Much?', *Journal of Urban Economics*, 110(C), 26–34, <https://doi.org/10.1016/j.jue.2019.01.003>.
- Britton, J., Erve, L. van der, Waltmann, B., and Xu, X. (2021), 'London Calling? Higher Education, Geographical Mobility and Early-Career Earnings', Department for Education Research Report, <https://doi.org/10.1920/re.ifs.2021.0198>.
- Card, D., Rothstein, J., and Yi, M. (2021), 'Location, Location, Location', US Census Bureau, Center for Economic Studies, Working Paper, <https://www2.census.gov/ces/wp/2021/CES-WP-21-32.pdf>.
- Carrascal-Incera, A., McCann, P., Ortega-Argilés, R., and Rodríguez-Pose, A. (2020), 'UK Interregional Inequality in a Historical and International Comparative Context', *National Institute Economic Review*, 253, R4–17, <https://doi.org/10.1017/nie.2020.26>.
- Carter, A. (2021), 'What Does the Public Think about Levelling Up?', Centre for Cities, blog post, <https://www.centreforcities.org/blog/what-does-the-public-think-about-levelling-up/>.
- Case, A., and Kraftman, L. (forthcoming), 'Health Inequalities', IFS Deaton Review of Inequalities.
- Cheshire, P., and Hilber, C. (2019), 'Housing', Centre for Economic Performance, CEP Election Analysis Paper 051.
- Chetty, R., and Hendren, N. (2018), 'The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects', *Quarterly Journal of Economics*, 133(3), 1107–62, <https://doi.org/10.1093/qje/qjy007>.
- Chetty, R., Hendren, N., and Katz, L. F. (2016), 'The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment', *American Economic Review*, 106(4), 855–902, <https://doi.org/10.1257/aer.20150572>.
- Chyn, E., and Katz, L. F. (2021), 'Neighborhoods Matter: Assessing the Evidence for Place Effects', *Journal of Economic Perspectives*, 35(4), 197–222, <https://doi.org/10.1257/jep.35.4.197>.
- Combes, P-P., Duranton, G., and Gobillon, L. (2008), 'Spatial Wage Disparities: Sorting Matters!', *Journal of Urban Economics*, 63(2), 723–42, <https://doi.org/10.1016/j.jue.2007.04.004>.
- Combes, P-P., and Gobillon, L. (2015), 'The Empirics of Agglomeration Economies', in G. Duranton, J. V. Henderson and W. C. Strange (eds), *Handbook of Regional and Urban Economics*, Elsevier, <https://doi.org/10.1016/B978-0-444-59517-1.00005-2>.
- Criscuolo, C., Martin, R., Overman, H. G., and Van Reenen, J. (2019), 'Some Causal Effects of an Industrial Policy', *American Economic Review*, 109(1), 48–85, <https://doi.org/10.1257/aer.20160034>.

Dauth, W., Findeisen, S., Moretti, E., and Suedekum, J. (2018), 'Matching in Cities', National Bureau of Economic Research, Working Paper 25227.

Davenport, A., and Zaranko, B. (2020), 'Levelling Up: Where and How?', in C. Emmerson, C. Farquharson and P. Johnson (eds), *The IFS Green Budget: October 2020*, <https://ifs.org.uk/publications/15055>.

De La Roca, J., and Puga, D. (2017), 'Learning by Working in Big Cities', *Review of Economic Studies*, 84(1), 106–42, <https://doi.org/10.1093/restud/rdw031>.

Duranton, G., and Puga, D. (2004), 'Micro-Foundations of Urban Agglomeration Economies', in J. Henderson and J-F. Thisse (eds), *Handbook of Regional and Urban Economics*, 4, Elsevier.

Duranton, G., and Puga, D. (2005), 'From Sectoral to Functional Urban Specialisation', *Journal of Urban Economics*, 57(2), 343–70, <https://doi.org/10.1016/j.jue.2004.12.002>.

Duranton, G., and Puga, D. (2020), 'The Economics of Urban Density', *Journal of Economic Perspectives*, 34(3), 3–26, <https://doi.org/10.1257/jep.34.3.3>.

Dustmann, C., Kastis, Y., and Preston, I. (forthcoming), 'Immigration and Inequality', IFS Deaton Review of Inequalities.

Fleury, N., and Gilles, F. (2018), 'The Intergenerational Transmission of Education: A Meta-Regression Analysis', *Education Economics*, 26(6), 557–73, <https://doi.org/10.1080/09645292.2018.1517863>.

Fujita, M., Krugman, P., and Venables, A. J. (1999), *The Spatial Economy: Cities, Regions, and International Trade*, MIT Press.

Gibbons, S., Overman, H., and Pelkonen, P. (2014), 'Area Disparities in Britain: Understanding the Contribution of People vs. Place through Variance Decompositions', *Oxford Bulletin of Economics and Statistics*, 76(5), 745–763, <https://doi.org/10.1111/obes.12043>.

Gibbons, S., Silva, O., and Weinhardt, F. (2013), 'Everybody Needs Good Neighbours? Evidence from Students' Outcomes in England', *Economic Journal*, 123(571), 831–74, <https://doi.org/10.1111/eoj.12025>.

Giupponi, G., and Machin, S. (forthcoming), 'Labour Market Inequality', IFS Deaton Review of Inequalities.

Glaeser, E., and Resseger, M. (2010), 'The Complementarity between Cities and Skills', *Journal of Regional Science*, 50(1), 221–44, <https://doi.org/10.1111/j.1467-9787.2009.00635.x>.

Gokan, T., Kichko, S., and Thisse, J-F. (2019), 'How Do Trade and Communication Costs Shape the Spatial Organization of Firms?', *Journal of Urban Economics*, 113, 103191, <https://doi.org/10.1016/j.jue.2019.103191>.

Goos, M., Manning, A., and Salomons, A. (2014), 'Explaining Job Polarization: Routine-Biased Technological Change and Offshoring', *American Economic Review*, 104(8), 2509–26, <https://doi.org/10.1257/aer.104.8.2509>.

Gyourko, J., Mayer, C., and Sinai, T. (2013), 'Superstar Cities', *American Economic Journal: Economic Policy*, 5(4), 167–99, <https://doi.org/10.1257/pol.5.4.167>.

Henderson, J., and Ono, Y. (2008), 'Where Do Manufacturing Firms Locate Their Headquarters?', *Journal of Urban Economics*, 63(2), 431–50, <https://doi.org/10.1016/j.jue.2007.02.006>.

Hilber, C. A. L., and Vermeulen, W. (2016), 'The Impact of Supply Constraints on House Prices in England', *Economic Journal*, 126(591), 358–405, <https://doi.org/10.1111/ecoj.12213>.

Jaffe, A. B., Trajtenberg, M., and Henderson, R. (1993), 'Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations', *Quarterly Journal of Economics*, 108(3), 577–98, <https://doi.org/10.2307/2118401>.

Joseph Rowntree Foundation (2021), 'Housing Costs', <https://www.jrf.org.uk/report/housing-costs-and-poverty-private-rents-compared-local-earnings>.

Joyce, R., Pope, T., and Roantree, B. (2019), 'The Characteristics and Incomes of the Top 1%', IFS Briefing Note 253, <https://ifs.org.uk/uploads/BN253-Characteristics-and-Incomes-Of-The-Top-1%25.pdf>.

Kerr, W. R., and Robert-Nicoud, F. (2020), 'Tech Clusters', *Journal of Economic Perspectives*, 34(3), 50–76, <https://doi.org/10.1257/jep.34.3.50>.

Knies, G., and Melo, P. C. (2019), 'Investigating People-Place Effects in the UK Using Linked Longitudinal Survey and Administrative Records', project report for Nuffield Foundation, <https://www.nuffieldfoundation.org/wp-content/uploads/2019/11/Investigating-people-place-effects-in-the-UK-final-report.pdf>.

Knies, G., Melo, P. C., and Zhang, M. (2021), 'Neighbourhood Deprivation, Life Satisfaction and Earnings: Comparative Analyses of Neighbourhood Effects at Bespoke Scales', *Urban Studies*, 58(13), 2640–59, <https://doi.org/10.1177/0042098020956930>.

Langella, M., and Manning, A. (2019), 'Residential Mobility and Unemployment in the UK', Centre for Economic Performance, Discussion Paper 1639, <https://cep.lse.ac.uk/pubs/download/dp1639.pdf>.

Manning, A., and Petrongolo, B. (2017), 'How Local Are Labor Markets? Evidence from a Spatial Job Search Model', *American Economic Review*, 107(10), 2877–907, <https://doi.org/10.1257/aer.20131026>.

McCann, P. (2019), 'Perceptions of Regional Inequality and the Geography of Discontent: Insights from the UK', *Regional Studies*, 54(2), 256–67, <https://doi.org/10.1080/00343404.2019.1619928>.

Moretti, E. (2004a), 'Human Capital Externalities in Cities', in J. V. Henderson and J-F. Thisse (eds), *Handbook of Regional and Urban Economics: Cities and Geography*, Elsevier.

Moretti, E. (2004b), 'Workers' Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions', *American Economic Review*, 94(3), 656–90, <https://doi.org/10.1257/0002828041464623>.

Moretti, E. (2013), 'Real Wage Inequality', *American Economic Journal: Applied Economics*, 5(1), 65–103, <https://doi.org/10.1257/app.5.1.65>.

Moretti, E. (2021), 'The Effect of High-Tech Clusters on the Productivity of Top Inventors', *American Economic Review*, 111(10), 3328–75, <https://doi.org/10.1257/aer.20191277>.

- Mulheirn, I., and Pisani, M. (2008), 'Working Tax Credit and Labour Supply', HM Treasury Economic Working Paper 3, https://revenuebenefits.org.uk/pdf/wtc_and_labour_supply_2008.pdf.
- Office for National Statistics (2020), 'Life Expectancy for Local Areas of the UK: between 2001 to 2003 and 2017 to 2019', <https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/healthandlifeexpectancies/bulletins/lifeexpectancyforlocalareasoftheuk/between2001to2003and2017to2019>.
- Oreopoulos, P. (2006), 'Estimating Average and Local Average Treatment Effects of Education when Compulsory Schooling Laws Really Matter', *American Economic Review*, 96(1), 152–75, <https://doi.org/10.1257/000282806776157641>.
- Ottaviano, G., and Thisse, J-F. (2004), 'Agglomeration and Economic Geography', in J. V. Henderson and J-F. Thisse (eds), *Handbook of Regional and Urban Economics: Cities and Geography*, Elsevier.
- Papoutsaki, D., Buzzeo, J., Gray, H., Williams, M., Cockett, J., Akehurst, G., Alexander, K., Newton, B., and Pollard, E. (2020), 'Moving Out to Move On', Institute for Employment Studies / Social Mobility Commission report, <http://www.employment-studies.co.uk/resource/moving-out-move>.
- Rae, A. and Nyanzu, E. (2019), 'An English Atlas of Inequality', <https://alasdairrae.github.io/atlasofinequality/>.
- Redding, S. J., and Turner, M. A. (2015), 'Transportation Costs and the Spatial Organization of Economic Activity', in G. Duranton, J. V. Henderson and W. C. Strange (eds), *Handbook of Regional and Urban Economics*, Elsevier.
- Rice, P., and Venables, A. (2021), 'The Persistent Consequences of Adverse Shocks: How the 1970s Shaped UK Regional Inequality', *Oxford Review of Economic Policy*, 37(1), 132–51, <https://doi.org/10.1093/oxrep/graa057>.
- Roback, J. (1982), 'Wages, Rents, and the Quality of Life', *Journal of Political Economy*, 90(6), 1257–78, <http://www.jstor.org/stable/1830947>.
- Rosen, S. (1974), 'Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition', *Journal of Political Economy*, 82(1), 34–55, <http://www.jstor.org/stable/1830899>.
- Rosen, S. (1981), 'The Economics of Superstars', *American Economic Review*, 71(5), 845–58, <https://www.jstor.org/stable/1803469>.
- Schlüter, T. (2013), 'Real Wages, Amenities and the Adjustment of Working Hours across Regional Labour Markets', Spatial Economics Research Centre, Discussion Paper SERCDP130, <http://eprints.lse.ac.uk/58529/>.
- Selby-Boothroyd, A. (2018), 'The Challenges of Charting Regional Inequality', <https://medium.economist.com/the-challenges-of-charting-regional-inequality-a9376718348>.
- University of Essex, Institute for Social and Economic Research. (2021). *Understanding Society: Waves 1-11, 2009-2020 and Harmonised BHPS: Waves 1-18, 1991-2009*. [data collection]. 14th Edition. UK Data Service. SN: 6614, DOI: 10.5255/UKDA-SN-6614-15.
- Weinhardt, F. (2014), 'Social Housing, Neighborhood Quality and Student Performance', *Journal of Urban Economics*, 82, 12–31, <https://doi.org/10.1016/j.jue.2014.06.001>.

Appendix A. Data

Data on wages come from the Annual Survey of Hours and Earnings (ASHE), a random 1% sample of employees in Great Britain which is regarded as the most accurate source of information on wages and earnings. We use data from 1998 to 2019. Wages are defined as the gross hourly wage, inclusive of bonuses and overtime. We assign people to places based on their local authority (LA) of work, as place of residence is not available in earlier years.

Data on employment come from the Labour Force Survey (LFS) and Annual Population Survey (APS). The LFS is a quarterly survey of the UK population with a sample size of around 85,000 per quarter. It provides the official measures of employment and unemployment. We pool the four quarters of the LFS in each year from 1998 to 2004. From 2004 to 2019, we use the APS, which combines the LFS with an additional boost sample to make it representative at the local authority level. It has a sample size of around 360,000 per year. We assign people to places based on their LA of residence.

Travel to work areas (TTWAs) are constructed as follows. We assign each LA to the TTWA with which it has the largest overlap in terms of postcodes. TTWAs are based on the 2011 Census; LAs are based on 2016 boundaries; and postcodes are based on the National Statistics Postcode Lookup in November 2016. We further group TTWAs where the sample size in either ASHE or APS is less than 200 in any year of the data with its nearest neighbour based on observed commuting flows. This leaves us with a total of 136 LA-based TTWAs.

Data on well-being are available in the APS from 2013 onwards. We use two common measures of well-being – happiness and life satisfaction – both scored on a 0–10 scale. As with employment, we assign people to places based on their LA of residence.

Data on gross value added (GVA) per capita, GVA per hour and household incomes (before and after housing costs) in Figure 5 come from the Office for National Statistics (ONS). GVA per capita uses the 'Regional Gross Value Added (Balanced) by Local Authority in the UK' data set.¹ GVA per hour uses the 'Subregional Productivity: Labour Productivity Indices by Local Authority District' data set.² Both are reported at the LA level, grouped into TTWAs as described above. Data on household incomes come from the 'Income Estimates for Small Areas, England and Wales' data set.³ Household incomes are modelled at the Middle Layer Super Output Area (MSOA) level using survey data from the Family Resources Survey, census data and administrative data, including on benefit claims and pay. We aggregate up to the (LA-based) TTWA level, weighting by the population size in each MSOA.

Earnings in Figure 5 are imputed as TTWA-level employment rate (APS) times TTWA-level average employee earnings (ASHE), as we do not have reliable data on self-employed earnings at the local level.

¹ <https://www.ons.gov.uk/economy/grossvalueaddedgva/datasets/regionalgrossvalueaddedbalancedbylocalauthorityintheuk>.

² <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/labourproductivity/datasets/subregionalproductivitylabourproductivityindicesbylocalauthoritydistrict>.

³ <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/datasets/smallareaincomeestimatesformiddlelayersuperoutputareasenglandandwales>.

Analysis by skill level uses the following classifications. For employment and well-being, we split individuals into graduates and non-graduates based on their highest qualification level. For wages – where information on degree status is unavailable – we use the classification in Aghion et al. (2019) to map four-digit Standard Occupational Codes to two skill groups: RQF 6 and above (which roughly corresponds to graduates) and RQF 5 and below (which roughly corresponds to non-graduates). In the full ASHE panel, which runs from 1998 to 2019, 14% of workers change occupations in a way that moves them between skill groups, and we classify them based on their modal skill level. Pooling across 1998–2019, 22% of individuals in the employment and well-being sample are graduates, and 27% of workers in the wage sample fall into the high-skilled group.

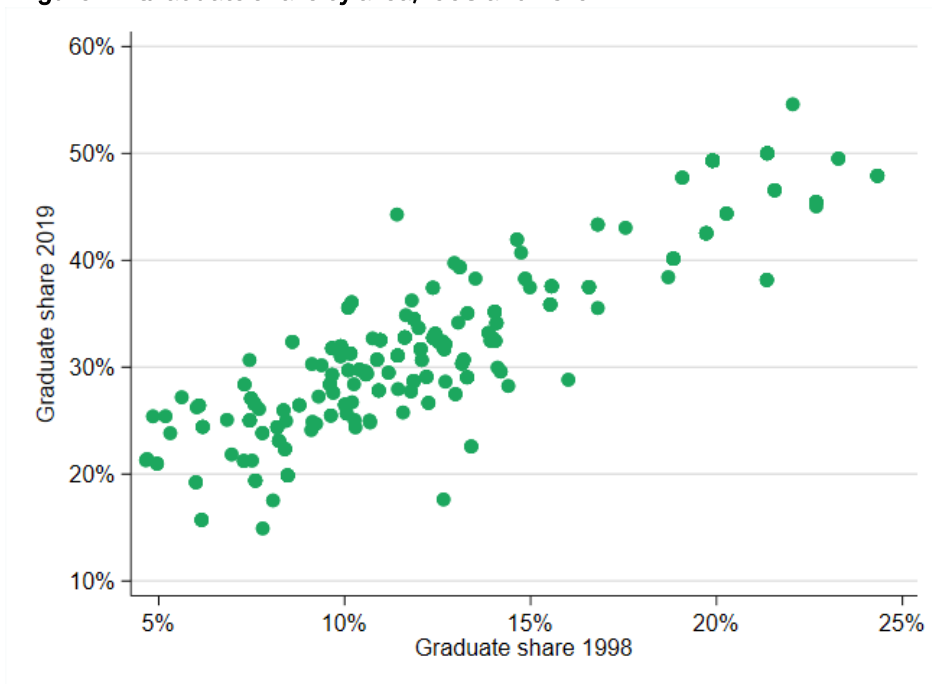
Data on rents in England come from the 'Private Rental Market Summary Statistics in England' data set, produced by the ONS.⁴ Data for Wales and Scotland come from their respective national statistics authorities.⁵ We use average rents for two-bedroom properties in our main analysis, and we use average rents for three-bedroom properties as a robustness check. LA averages are aggregated to (LA-based) TTWAs weighted by population size.

⁴ <https://www.ons.gov.uk/peoplepopulationandcommunity/housing/datasets/privaterentalmarketsummarystatisticsinengland>.

⁵ <https://gov.wales/private-sector-rents> for Wales; <https://www.gov.scot/publications/private-sector-rent-statistics-2010-2019/> for Scotland.

Appendix B. Additional figures

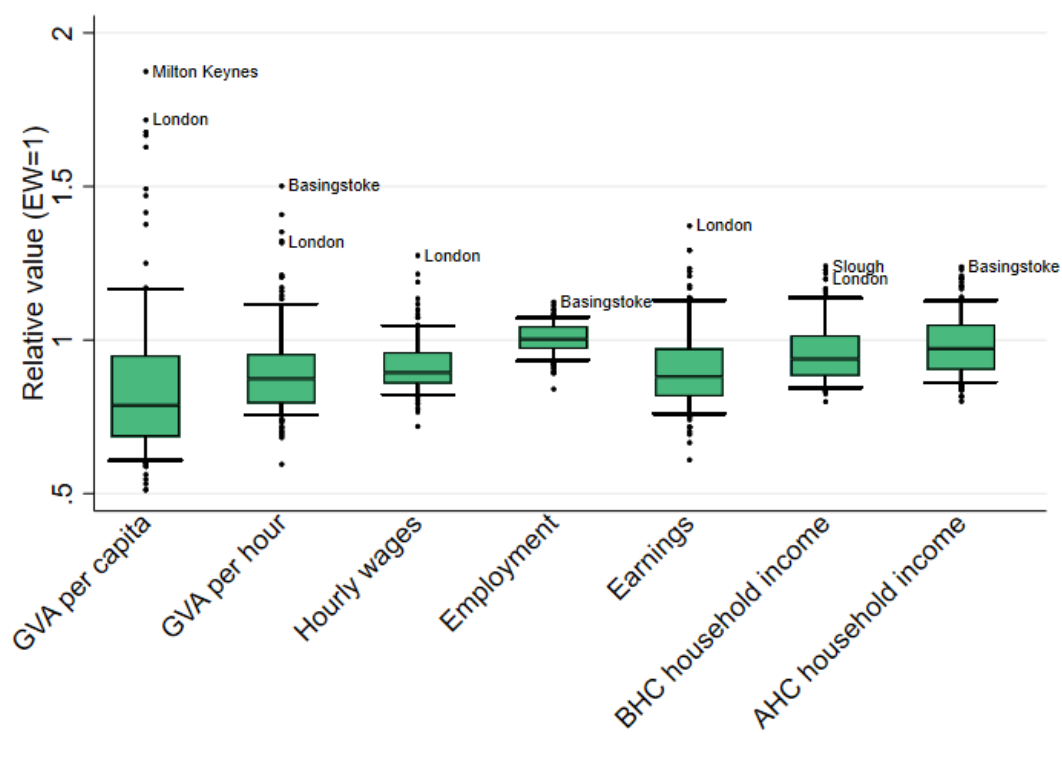
Figure B1. Graduate share by area, 1998 and 2019



Note: Working-age (16–64) population, excluding those in full-time education. Figure plots TTWA graduate shares in 1998 against graduate shares in 2019 and shows that the share has risen across all areas. The expansion of higher education resulted in larger percentage increases in graduate shares in areas with initially lower shares of graduates.

Source: ONS.

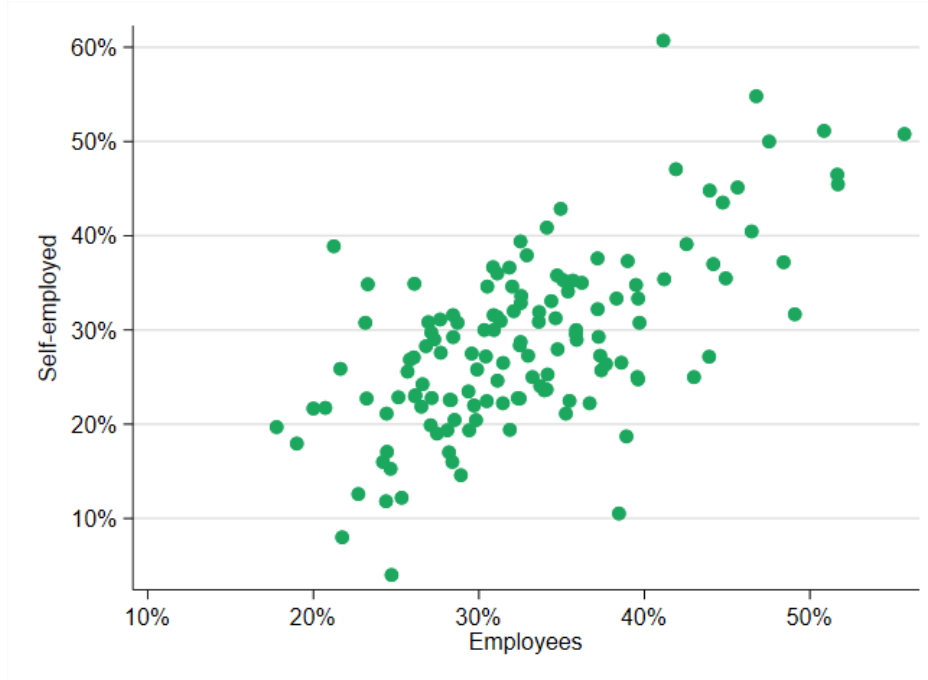
Figure B2. Distribution of labour market outcomes across areas, 2018, England and Wales only



Note: Figure replicates Figure 5 using data for England and Wales (EW) only. The first five plots in Figure 5 are drawn for Great Britain, while the two plots for household income before and after housing costs are for England only, due to data availability. The distribution of outcomes in Figure B2, using data for England and Wales only, looks very similar to that in Figure 5. For further details, see notes to Figure 5 and the main text.

Source: ONS (GVA per capita, GVA per hour and household income); APS (employment and earnings); ASHE (hourly wages and earnings).

Figure B3. Share of employees and self-employed workers with degrees, 2018



Note: Working-age (16–64) population, excluding those in full-time education. Figure shows that the share of self-employed workers with a degree by area is highly correlated with the share of employees with a degree, although the degree of variation for the share of self-employed is slightly higher. If this translates into larger variation in wages, then our data understate the dispersion of wages including the self-employed. However, the high correlation suggests our main findings would continue to hold.

Source: APS.

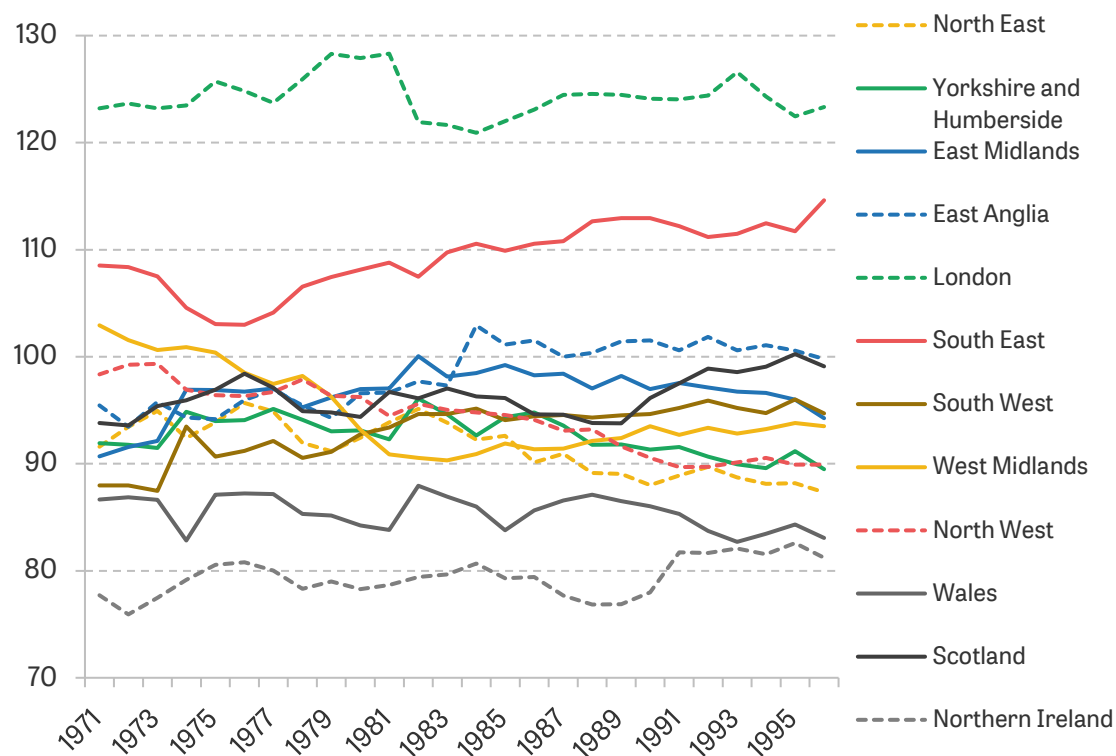
Figure B4. Correlation between employment shocks in 1970s and subsequent male employment rates



Note: Panel a shows that unemployment shocks in the 1970s – as measured by the change in male employment rates (MERs) between 1971 and 1981 – were highly spatially correlated and that there was high persistence until at least 2011 as evidenced by the correlation between changes 1971–81 and changes 1971–2011. Panel b provides further evidence for this persistence as evidenced by the lack of correlation between changes 1981–2011 and changes 1971–81. Places that had a big initial shock did not see much recovery in the next three decades.

Source: Rice and Venables, 2021.

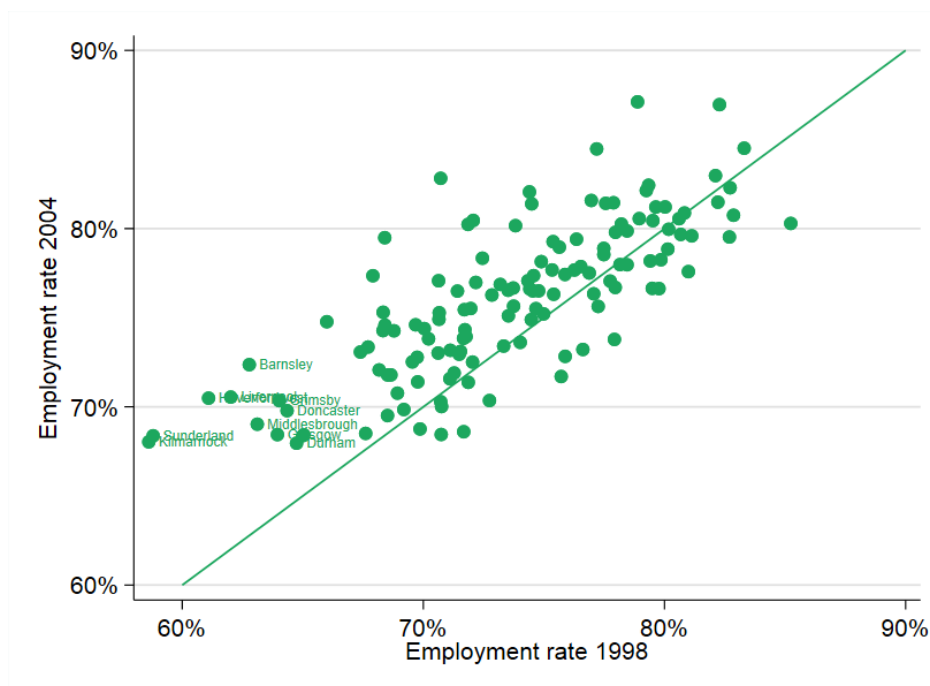
Figure B5. Relative GDP per capita by region (UK=100), 1971–96



Note: Figure shows that GDP per capita in the West Midlands and the North West fell between 1971 and 1996 relative to the UK average, whilst GDP per capita in the South East and the South West rose. There was little change in disparities across the other government standard regions over this period.

Source: ONS.

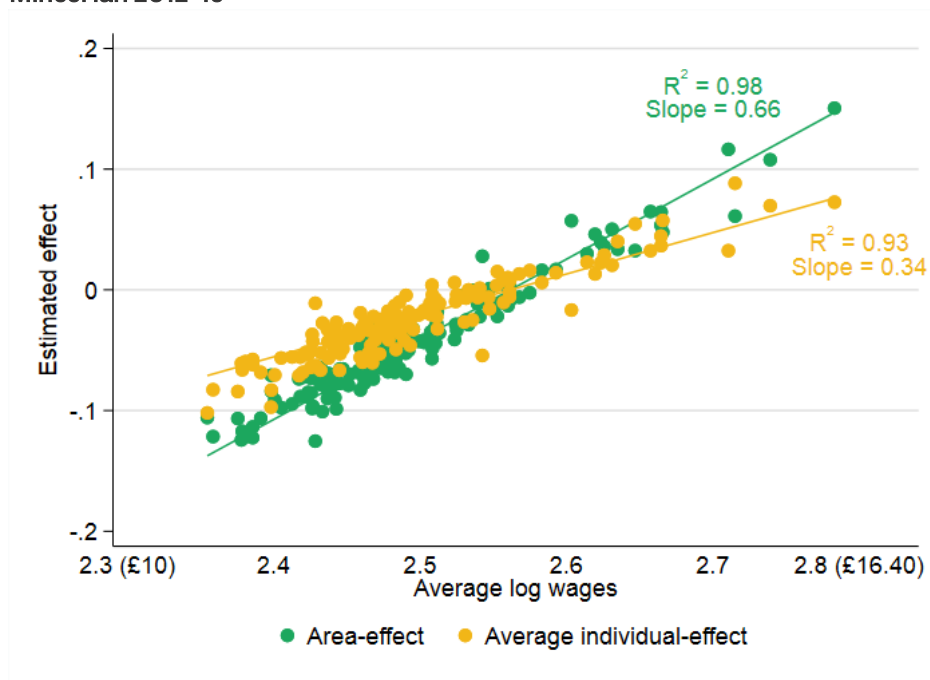
Figure B6. Employment rate by area, 1998 and 2004



Note: Figure shows that areas with very low employment rates in 1998 saw large increases in employment between 1998 and 2004.

Source: LFS.

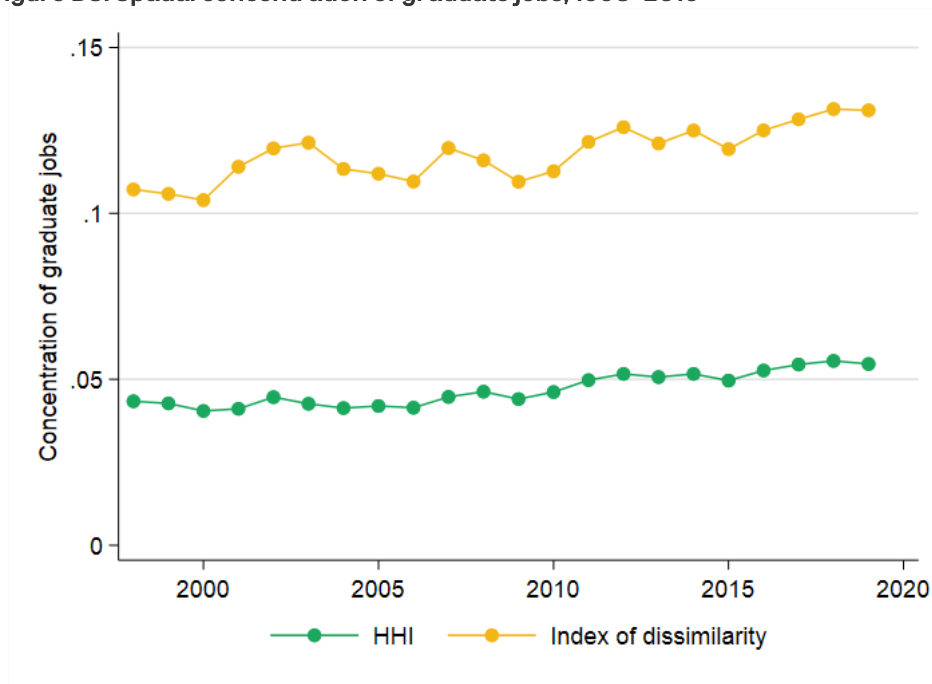
Figure B7. Correlation between estimated area effects and average person effects in wages, Mincerian 2012–19



Note: Working-age (16–64) population. Figure replicates Figure 14(a) for wages using a Mincerian regression. It shows that the relative slopes of the two lines are now more comparable to those in Figure 14(b) for employment, which is estimated using a Mincerian regression.

Source: ASHE.

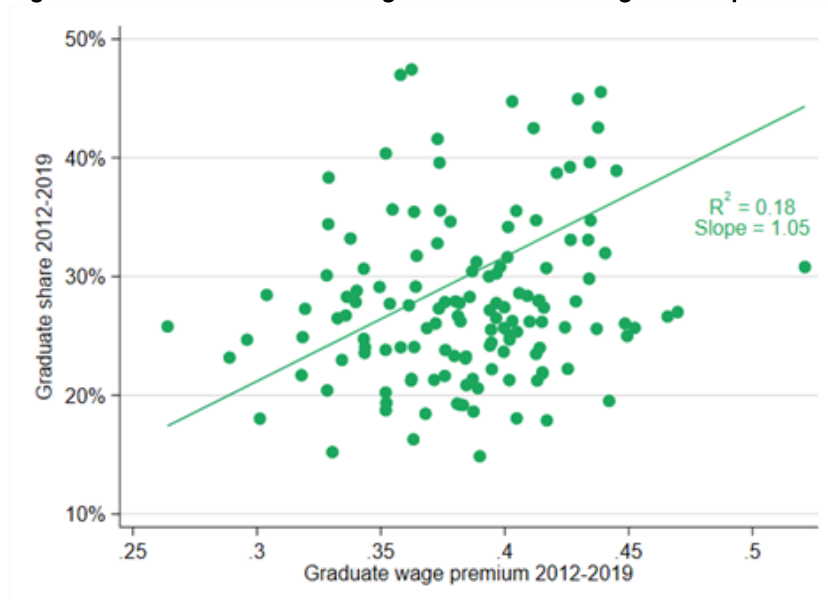
Figure B8. Spatial concentration of graduate jobs, 1998–2019



Note: Working-age (16–64) population. Figure shows that measures of the spatial concentration of graduate jobs, defined using the Aghion et al. (2019) classification, increased steadily from 1998 to 2019.

Source: ASHE.

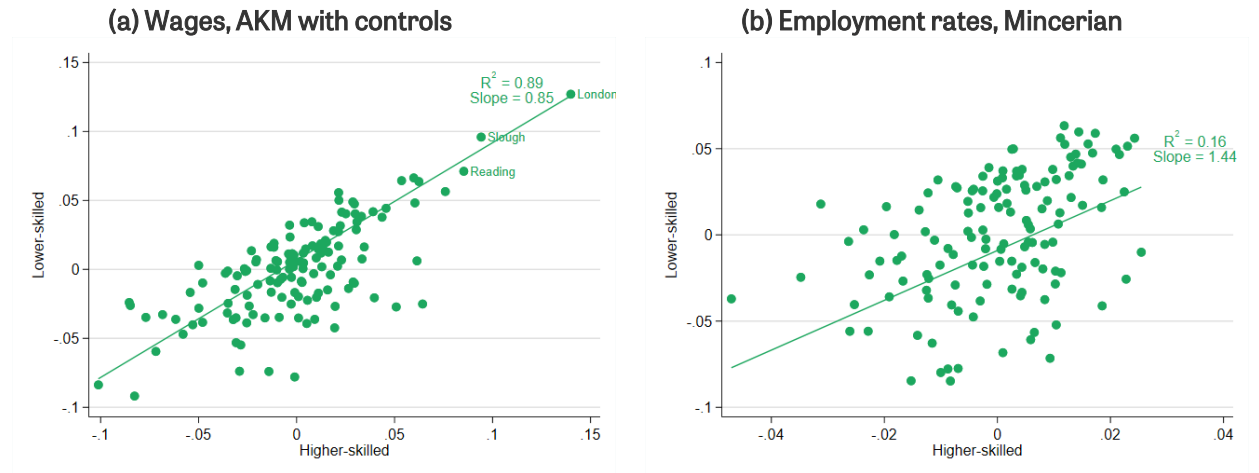
Figure B9. Correlation between graduate share and graduate premium, 2012–19



Note: Working-age (16–64) population. Graduate wage premiums estimated using regression of log wages on degree status interacted with TTWA dummies, sex, five-year age group and year dummies. Regression line weighted by TTWA-level population size.

Source: APS.

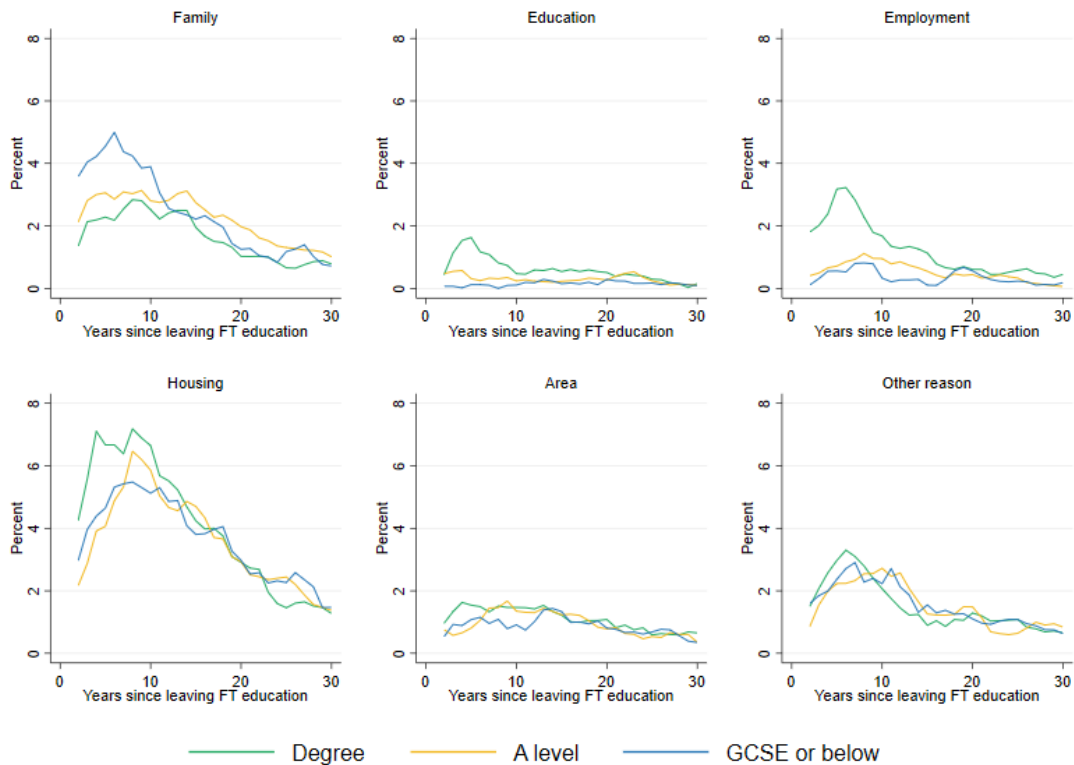
Figure B10. Correlation between area effects for higher- and lower-skilled workers



Note: Working-age (16–64) population. Figure shows that there is a high correlation between estimated area effects for higher- and lower-skilled groups, especially for wages. 'Higher-skilled' refers to RQF 6+ in panel a and graduates in panel b. Area effects for wages estimated using 'AKM with controls' specification. Area effects for employment estimated using Mincerian regression. Regression line weighted by TTWA-level population size.

Source: APS; ASHE.

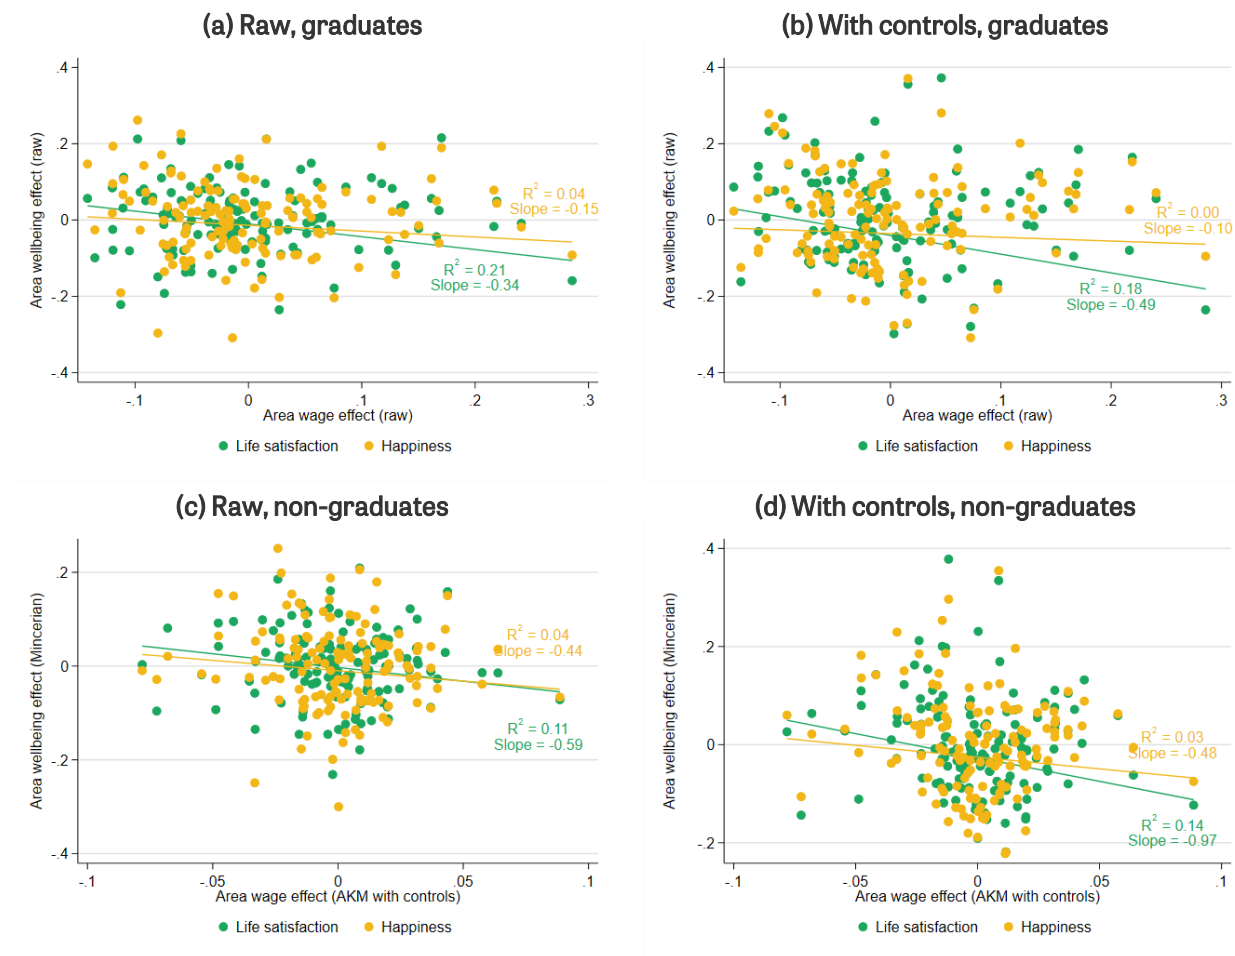
Figure B11. Share of individuals who move TTWA in any given year, by reason and education, 2009–19



Note: Working-age (16–64) population, excluding those in full-time education. Figure shows that lower-skilled individuals are less likely to move for employment reasons, and more likely to move for family reasons, than those with degrees.

Source: Understanding Society waves 1–10.

Figure B12. Correlation between area-level wages and well-being, by degree status, 2012/2013 to 2019



Note: Working-age (16–64) population. Uses data on wages from 2012 to 2019 and data on well-being from 2013 to 2019. Figure shows that there is no positive correlation between wages and well-being, for both graduates and non-graduates, even though non-graduates are much less mobile than graduates. Regression line weighted by TTWA-level population size.

Source: APS; ASHE.

Appendix C. Details on estimating equations and variance decomposition

Estimating equations

We follow the urban economics literature – see Card, Rothstein and Yi (2021) for a recent application and Combes and Gobillon (2015) for a review – and estimate area effects for wages as follows. First, we consider raw differences in (log) wages across areas of work d_r , ignoring differences in individual characteristics. We pool the post-recession years of data (2012–19) and include year fixed effects λ_t :

$$(1) \quad \text{Raw: } \ln w_{irt} = d_r + \lambda_t + \varepsilon_{irt}$$

Next, we run a Mincerian regression of wages on area dummies d_r and observable individual characteristics x_{it} : gender, age, skill level and whether they work full-time or part-time. Education is not included in the ASHE data, so we use a measure of skill derived from occupations (four-digit SOC code) following Aghion et al. (2019). The occupation match comes from the current UK immigration rules, and we split people into two skill groups, high-skilled (RQF 6 and above, or roughly graduates) and low-skilled (below RQF 6, or roughly non-graduates).

$$(2) \quad \text{Mincerian: } \ln w_{irt} = d_r + x'_{it}\beta + \lambda_t + \varepsilon_{irt}$$

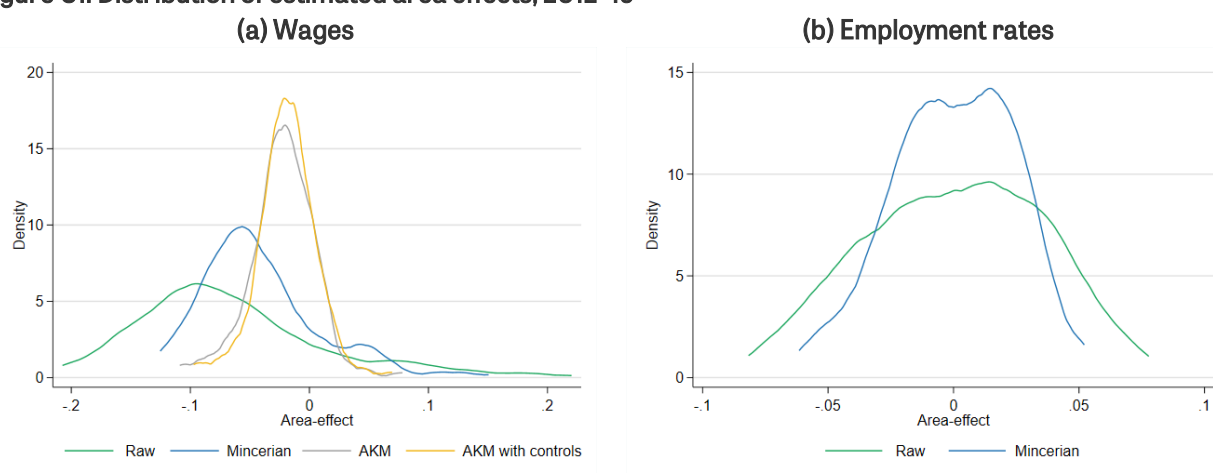
The Mincerian regression above does not control for unobservable characteristics. To deal with this, we follow the same individuals as they move across areas in a two-way fixed effects (AKM) framework. We run two specifications, one with only individual fixed effects α_i , and one including age and full-time status. These regressions use data on both movers and non-movers: area effects are identified from movers, while all workers are used to estimate the effect of time-variant characteristics.

$$(3) \quad \text{AKM: } \ln w_{irt} = d_r + \alpha_i + \lambda_t + \varepsilon_{irt}$$

$$(4) \quad \text{AKM with controls: } \ln w_{irt} = d_r + \alpha_i + age_{it} + fulltime_{it} + \lambda_t + \varepsilon_{irt}$$

As noted by Combes, Duranton and Gobillon (2008) and recognised in the subsequent literature, identification in the AKM framework requires that moves across areas are as good as random, conditional on observables and time-varying unobservables. If mobility is correlated with shocks to wages, or if the choice of area is partly driven by idiosyncratic match effects, the individual error term (ε_{irt}) may be correlated with d_{rt} . Card, Rothstein and Yi (2021) present several tests that support the random mobility assumption in the context of US local labour markets.

We implement an analogous approach to assess the effect of area and individual characteristics on employment but drop the individual fixed effects specifications as we do not have panel data for employment. The Mincerian regression includes controls for gender, age, education, ethnicity, whether UK-born, whether a UK citizen and household characteristics interacted with gender (marital status, number of children and age of the youngest child).

Figure C1. Distribution of estimated area effects, 2012–19

Note: Working-age (16–64) population only. Area effects normalised to average across TTWAs.

Source: APS; ASHE.

Table 1 in the main text uses results from estimating these different specifications to provide a snapshot of area effects at different points of the distribution. Figure C1 plots the entire distribution for each of the four different specifications for wages and two for employment. The effect of area differences in individual characteristics is clearly visible. Controlling for observable characteristics reduces the spread of area effects and controlling for unobservable individual fixed effects (in wages) reduces the spread even further.

Variance decomposition

Gibbons, Overman and Pelkonen (2014) propose several decompositions to disentangle the contribution of area and individual effects to overall wage differences and to area differences in wages. We consider one of these: the correlated variance share. Since we are interested in understanding spatial disparities (as opposed to inequality between individuals), we mainly apply this to differences in area-level means rather than differences in individual outcomes.

As an example, consider the 'AKM with controls' regression on wages above. The variance in average wages across areas can be expressed as $var(\ln \bar{w}_r) = var(\hat{d}_r + \bar{\alpha}_r + \hat{\beta} \bar{X}_r + \bar{\varepsilon}_r)$, where X_r denotes observable controls $age_{it} + fulltime_{it} + \lambda_t$. This can be decomposed as

$$(5) \quad var(\ln \bar{w}_r) = var(\hat{d}_r) + var(\bar{\alpha}_r) + var(\hat{\beta} \bar{X}_r) + 2cov(\hat{d}_r, \bar{\alpha}_r) + 2cov(\hat{d}_r, \hat{\beta} \bar{X}_r) + 2cov(\bar{\alpha}_r, \hat{\beta} \bar{X}_r) + var(\bar{\varepsilon}_r)$$

The correlated variance share is defined as $var(\hat{d}_r)/var(\ln \bar{w}_r)$ – the share of the variance in area-level mean wages that is accounted for by the variance in estimated area effects.

Card, Rothstein and Yi (2021) use an alternative approach that decomposes the variance in raw mean wages across areas using the fact that the variance can be expressed as a weighted average of the variance in area effects and the variance in area-level average individual effects. The weight on each component is equal to the slope of a univariate regression of that component on the area-level mean wage, since $var(\ln \bar{w}_r) = cov(\ln \bar{w}_r, \hat{d}_r + \bar{\alpha}_r + \bar{X}_{it}\hat{\beta})$ and the slope of a regression of \hat{d}_r on $\ln \bar{w}_r$ is $cov(\ln \bar{w}_r, \hat{d}_r)/var(\hat{d}_r)$. As discussed in Box 2 in the main text, this comparison does not separate out the correlation between individual and area effects that forms

the third part of the correlated variance share decomposition – $2Cov(\hat{d}_r, \bar{\alpha}_r) + 2Cov(\hat{d}_r, \hat{\beta}\bar{X}_r)$ – used in the main text.

Table C1 uses the correlated variance share to decompose the variance in area-level mean wages and employment rates in 2012–19 into area effects, individual effects (individual fixed effects and time-varying observables), and correlations between these components. The columns correspond to the four regression equations set out above. The first row shows the variance in area-level means. The following rows show the share of the variance in area-level means that is accounted for by each of the components.

Although we focus on decomposing the variance in area-level average wages and employment rates, the same variance share can be used to calculate the share of individual inequality accounted for by area effects. Using the notation above, this is $var(\hat{d}_r)/var(\ln w_{irt})$, with \hat{d}_r estimated using equations 1 to 4. For wages, the share of individual inequality accounted for by area effects is 5.3% using raw differences (equation 1), 2.4% using the Mincerian regression (equation 2), 0.6% using the AKM regression (equation 3) and 0.5% using the 'AKM with controls' regression (equation 4). For employment, the corresponding shares are 0.6% using raw differences and 0.3% using the Mincerian regression.

Table C1. Variance decomposition, 2012–19

	Wages			Employment
	Mincerian	AKM	AKM controls	Mincerian
Variance	0.0064	0.0064	0.0064	0.0013
Area effects	40.03	12.22	9.69	46.23
Individual fixed effect		58.75	52.42	
Individual observables	18.22	0.10	9.26	19.79
Correlations (2 x covariance)				
Area effects, individual FE		27.12	17.24	
Area effects, individual observables	42.81	0.19	8.70	33.98
Individual FE, individual observables		0.72	2.14	

Note: Working-age (16–64) population. Columns refer to regression equations 1 to 4 above. Rows show percentage of variance accounted for by each component.

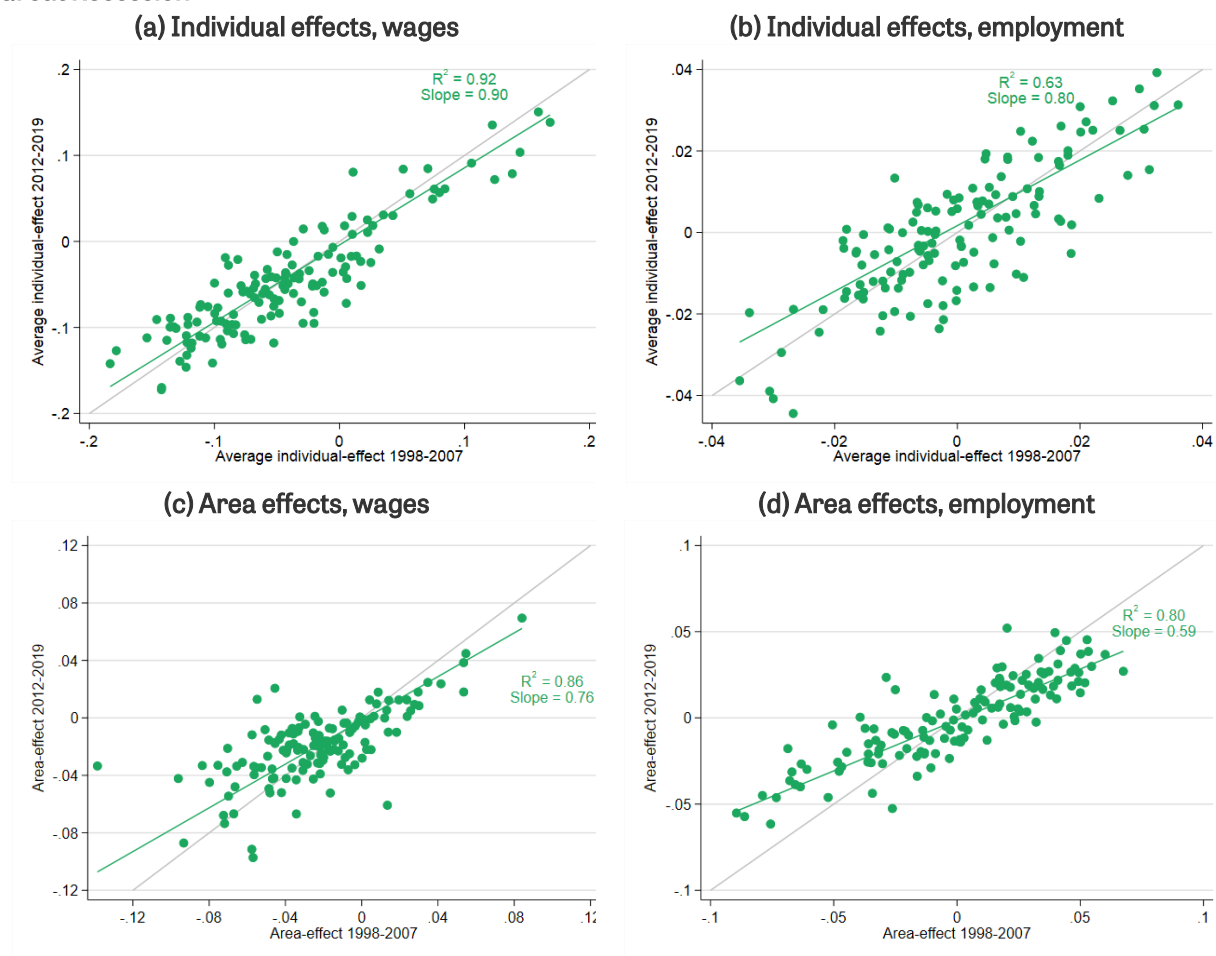
Source: APS; ASHE.

Changes over time

When looking at changes over time, we see persistence in average individual and area effects, with the persistence for wages higher than for employment. Figure C2 shows that places' relative positions – in terms of both the composition of workers and the area effect on wages and employment probabilities – do not move much between the periods before and after the Great Recession (1998–2007 and 2012–19 respectively). Figure C3 – which plots the distribution of average individual effects and area effects in the two different periods – shows less convergence in individual effects than in area effects. This suggests that the extent to which spatial disparities are explained by differences across areas in individual characteristics may have increased over time.

Table C2 shows the decomposition of area-level means in wages and employment in the pre-Great-Recession period. Comparing Tables C1 and C2, we see that area effects accounted for 13% of the variation in raw mean wages across areas in 1998–2007 and 10% of the variation in 2012–19. That is, area effects account for a smaller share of area disparities in the post-recession period, with area differences in individual characteristics playing a bigger role. Area effects accounted for 67% of the variation in employment rates across areas in 1998–2007 and only 46% of the variation in 2012–19.

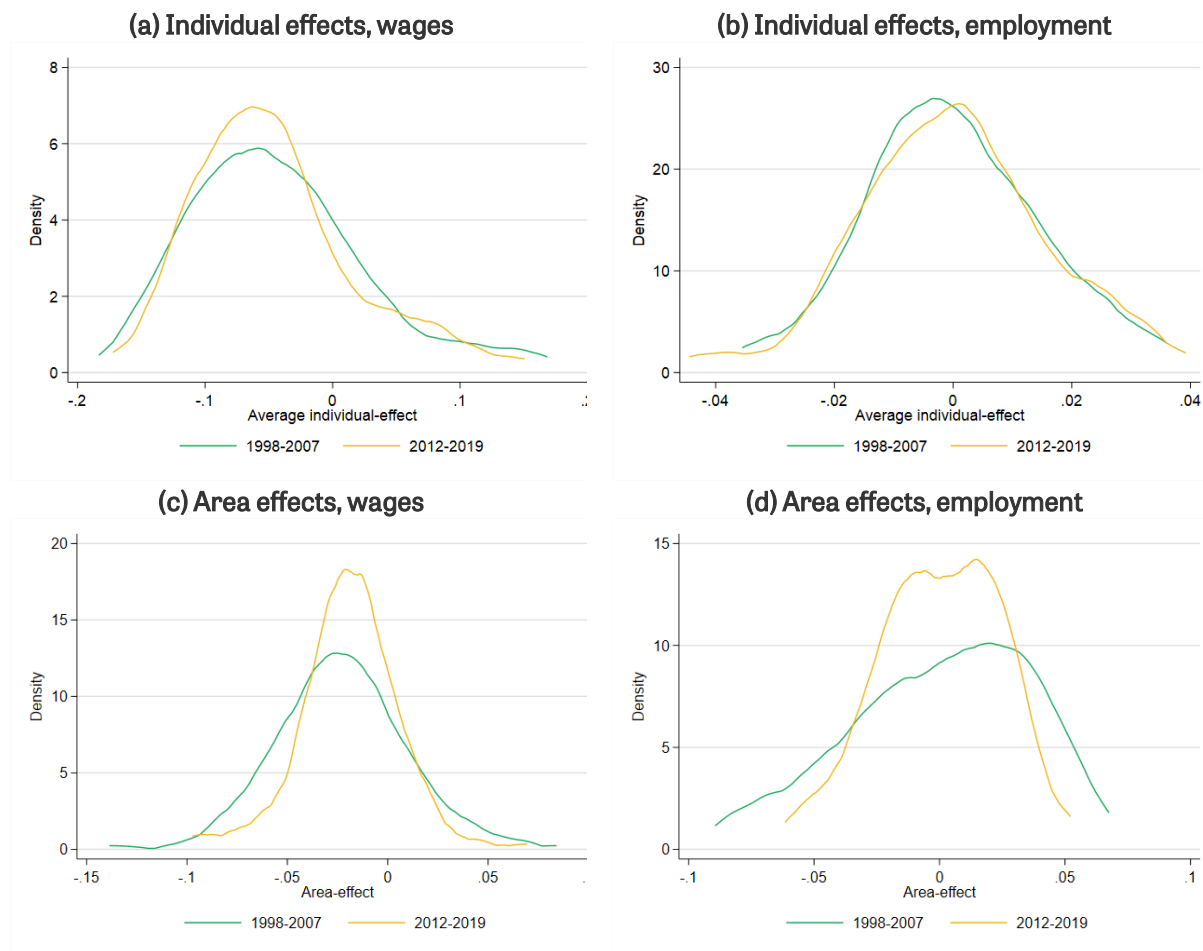
Figure C2. Persistence in individual and area effects in wages and employment, before and after Great Recession



Note: Working-age (16–64) population. Area and individual effects based on 'AKM with controls' regression for wages and 'Mincerian' regression for employment. Regression line weighted by TTWA-level population size.

Source: ASHE.

Figure C3. Distribution of individual and area effects in wages and employment, before and after Great Recession



Note: Working-age (16–64) population. Area and individual effects based on 'AKM with controls' regression for wages and 'Mincerian' regression for employment. Regression line weighted by TTWA-level population size.

Source: ASHE.

Table C2. Variance decomposition, 1998–2007

	Wages			Employment
	Mincerian	AKM	AKM controls	Mincerian
Variance	0.0085	0.0085	0.0085	0.0019
Area effects	51.22	14.66	12.84	66.64
Individual fixed effect		55.50	54.05	
Individual observables	10.99	0.57	1.73	11.30
Correlations (2 x covariance)				
Area effects, individual FE		28.92	25.90	
Area effects, individual observables	38.07	-0.25	2.78	22.06
Individual FE, individual observables		0.13	2.46	

Note: Working-age (16–64) population. Columns refer to regression equations 1 to 4 above. Rows show percentage of variance accounted for by each component.

Source: LFS; APS; ASHE.