Understanding the ethnicity pay gap in the London Borough of Ealing and how it can be addressed

The University of West London and Ealing Council

CREATING GOOD JOBS
TACKLING THE CLIMATE CRISIS
FIGHTING INEQUALITY

www.ealing.gov.uk

Contents

Acknowledgements

A special thank you to:

Vanessa Guthrie, the chairperson of Ealing's Race Equality Commission, Citizen Tribunal Income and Employment Subgroup, who inspired this work to happen.

Mark Coulson, Professor of Psychology at the School of Human and Social Sciences, for leading the research and authoring the report with rigour, clarity, and commitment.

Kasem Abbas, Senior Regeneration, Investment and Jobs Officer at Ealing Council, who answered questions, provided data, and engaged in endless discussions and requests. It was a pleasure working with you, Kas.

Sophie Beagles, Head of Regeneration, Investment and Jobs at Ealing Council, for timely and wise guidance.

Connor McDonagh, Assistant Director of Regeneration, Economy & Skills at Ealing Council, for providing support and advocacy.

Darren Bestley and Gurdeep Singh from the GIS team at Ealing Council for the provision of critical data.

Dr Jen Mayer, head of the School of Human and Social Sciences at the University of West London, for her support and advocacy of this work.

John Charlton, head of knowledge exchange at the University of West London, for enthusiastic help and support throughout the project.

Professors Tom Dickins and Ben Hine at the University of West London, for critical reading, sensemaking, and statistical minutiae.

Councillor Kamaljit Kaur Nagpal, Cabinet member for decent living incomes

Foreword

In the London Borough of Ealing (LBE), we are proud of our diversity. Over half our residents are from ethnic minorities and over 160 languages are spoken in the borough. But sadly, for too many of our residents, their ethnicity – directly or indirectly – acts as a barrier to prosperity and equal opportunities.

As set out in our council plan, we are committed to ensuring all our residents can access a decent living income. However, as is the case across the country, this report shows that an ethnicity pay gap exists in the borough, meaning our residents from ethnic minorities earn less than White workers.

Currently, ethnicity pay reporting is voluntary. The new Equality (Race and Disability) Bill will make this reporting compulsory for employers with at least 250 employees, however, currently there is lack of publicly available data and insight on this topic. Equally, this Bill will not require micro, small and medium businesses that make up 99.8% of the local economy to report their data. Therefore, we have partnered with the University of West London to fully understand this challenge in the borough, working closely with Ealing's Race Equality Commission.

We are the first west London borough to publish this type of report and we are already taking positive steps to address the findings. Our new jobs and skills strategy sets out how we will create more inclusive pathways for our residents to access employment in our growth sectors through our Good for Ealing, Learn Ealing, and Work Ealing programmes. Our work as the first London Living Wage Place in west London has seen nearly 1,000 low-paid workers in the borough receive pay rises, benefitting workers from the global majority who are over-represented in these low paid roles.

However, there is much more that needs to be done to ensure pay is equitable for all ethnic groups, and that opportunities to enter and progress in the labour market are truly equitable. This cannot be done by Ealing Council alone. It requires action from a range of stakeholders, including educational institutions, training providers, business groups, and employers across the borough. This research is intended to be a shared evidence base that can be used by these different organisations to take action.

To help us grow this evidence base in the future, we strongly encourage more organisations to report on their ethnicity pay data. Employees want more transparency on pay from their employers and it's a great way to demonstrate your commitment to equality and diversity. Being a more inclusive company can help attract talent, retain staff and improve employee performance. Understanding and tackling the ethnicity pay gap is good for our residents and good for our businesses.

Why are we doing this, and what do we hope to achieve?

This work stems directly from commitments outlined in Ealing's council plan and the mandate of the borough's Race Equality Commission (REC). Both identify tackling inequalities, including the ethnicity pay gap, as a key priority to ensure that all residents have access to equal opportunities and decent living incomes.

We need a clear understanding of the ethnicity pay gap challenge across our seven towns and the factors contributing to this. However, there is currently and will continue to be a lack of publicly available data and insight on this topic. Ethnicity pay gap reporting remains voluntary at this point in time. The new Equality (Race and Disability) Bill will make this reporting compulsory, but only for businesses with at least 250 employees, which means it won't apply to over 99% of businesses in the borough.

There is already good work underway, with **nearly 1,400 residents from ethnic minorities completing qualifications and training through our Learn Ealing programme** over the past two academic years. But we must want to go further. This research will inform the delivery of future initiatives to address the ethnicity pay gap across the borough, ensuring finite resources are effectively prioritised.

The recommendations outlined in this report are **short-term actions that the council, its partners, and Ealing's business community can undertake to drive forward positive change** and showcase Ealing as a borough committed to fairness and equity.

> *We need a clear understanding of the ethnicity pay gap challenge across our seven towns and the factors contributing to this.*

Executive summary

A national figure suggests that for every £1 earned by white employees, those from ethnic minorities earned 85.2p. This report explores this issue within the London Borough of Ealing specifically. Ealing is ethnically diverse, and contains the largest Polish, Afghan, and Serbian populations in England and Wales, the second largest Japanese and Iranian populations, and the fourth highest Arab population.

In England and Wales, the ethnicity pay gap (EPG) remained relatively unchanged between 2014-2019. The median hourly rate of pay increased for all ethnic groups over this period, but the pay gap between *White British/ Irish* workers and those from other ethnicities remains more or less unchanged. The median hourly rate of pay in 2019 for *White British/ Irish* workers was £15.71 compared to a range of £9.25 to £12.72 for workers from all other ethnicities.

Figure 1.
Median pay by ethnicity: England and Wales (excludes London) 2014 -19. ONS data

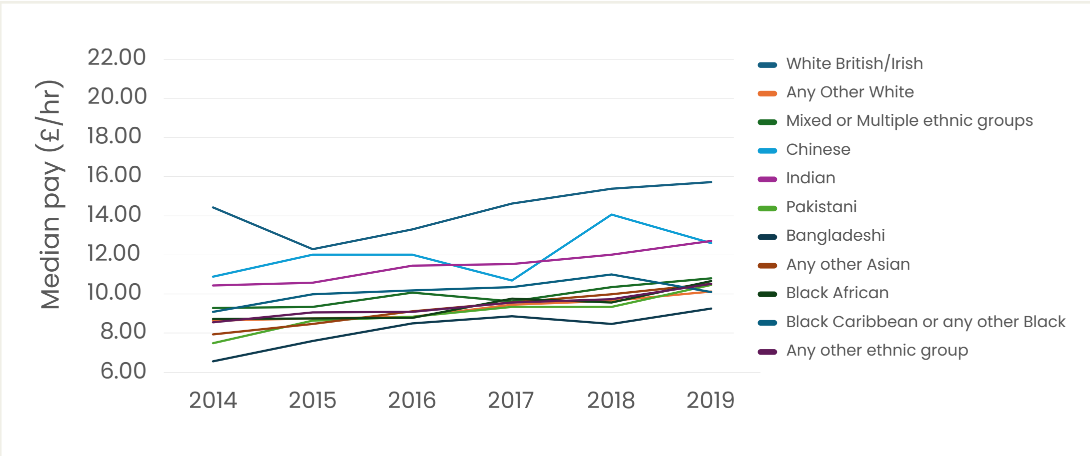

Ealing is ethnically diverse, and contains the largest Polish, Afghan, and Serbian populations in England and Wales, the second largest Japanese and Iranian populatins, and the fourth highest Arab population.

While earnings in London are higher than in the remainder of England and Wales, the EPG is more marked in the capital than elsewhere. The median hourly rate of pay in 2019 for *White British / Irish workers* in London was £21.63 compared to a range of £12.12 to £17.71 for workers from all other ethnicities.

Figure 2. Median pay by ethnicity: London only 2014 - 19. ONS data

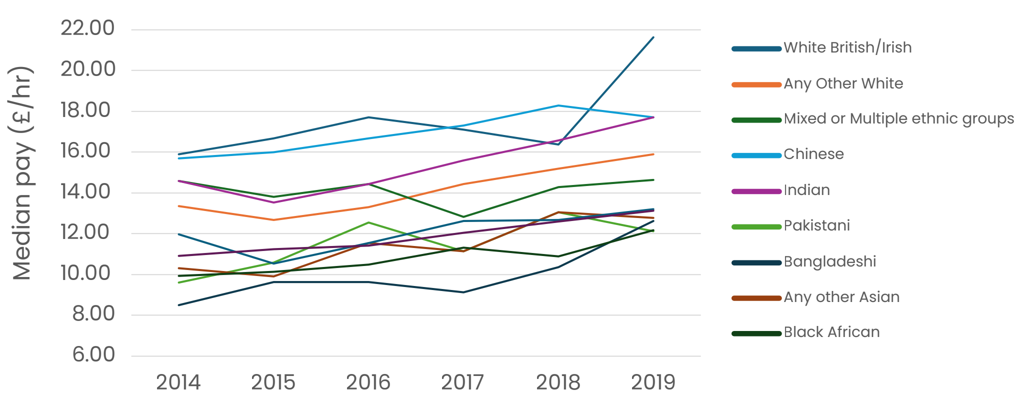

Ealing is a diverse borough, with variation across the seven towns.
In 2021, approximately 43% of residents identified as *White*, compared with 49% in 2011. Nearly a third identified as *Asian, Asian British, or Asian Welsh.*

In Ealing, income deprivation is greater when there are more *Black British African* or *Asian British Indian* people living in an area, with the number of *Black British African* people accounting for a **16% variation in income deprivation** and the number of *Asian British Indian* people accounting for a **9% variation in income deprivation**. That is, the greater the number of people from these groups who live in a Lower Layer Super Output Area (LSOA), the greater the deprivation.

In Ealing, eight ethnic populations are associated with lower household income –

- *Asian British Bangladeshi*
- *Asian British Indian*
- *Asian British Pakistani*
- *Asian British Other*
- *Black British African*
- *Black British Caribbean*
- *Black British Other*
- *Mixed White Black Caribbean*

That is, the greater the number of people from these groups who live in a Middle Layer Super Output Area (MSOA), the lower the average household income is in that MSOA.

In Ealing, three ethnic populations are associated with higher household income - *Mixed White Asian, White English Welsh Scottish,* or *White Irish.* That is, the greater the number of people from these groups who live in an MSOA, the greater the average household income is in that MSOA.

In Ealing, employment deprivation is greater when there are more Black British African people living in an area, with a variation of approximately 12%. That is, the greater the number of people from these groups who live in an LSOA, the greater the deprivation. No other ethnic groups are associated with increased employment deprivation.

In Ealing, education deprivation is greater when there are more Black British African people living in an area, with a variation of approximately 12%. That is, the greater the number of people from these groups who live in an LSOA, the greater the deprivation. No other ethnic groups are associated with increased education deprivation.

Higher house prices are associated with a higher number of *White English Welsh Scottish* living in the LSOA. That is, the size of the White English Scottish Welsh population is associated with higher average house process. No other ethnic groups are associated with house prices.

The new Equality (Race and Disability) Bill will make ethnicity pay reporting compulsory for employers with at least 250 employees and improve access to data to further bolster the findings of this report. However, more can be done to support businesses to analyse, report and address their EPG voluntarily. Our recommendations include the following and build on those already set out under Ealing's Jobs and Skills Strategy:

- **In preparation for the Bill, ensure large employers are ready to meet this statutory requirement and encourage them to publish reports in advance.** This could require targeted engagement, guidance and specialist business support.
- **Encourage employers of less than 250 employees to analyse and report on their EPG voluntarily.** This is particularly relevant in Ealing, where micro, small and medium businesses make up 99.8% of the local economy. This could require a communication campaign, guidance, and specialist business support.
- **Support businesses of all sizes to develop actions plans to address their EPG.** This could require online guidance and specialist business support.
- **Through Ealing's London Living Wage,** make targeted efforts to increase the number of London Living Wage accredited businesses in parts of the borough where there are more Black British African, or Asian British Indian people living and working, and income deprivation is greater.
- **Through Ealing's new growth sector forums,** explore how the EPG varies across the foundational and growth sectors that make-up Ealing's economy, and measures that need to be put in place to address this.

- **Through the council's purchasing powers,** encourage more businesses to report on and develop action plans to address the EPG. This should be a key consideration of the social value component of a procurement.
- **Encourage all commercial tenants of council property to report their EPG and develop action plans to address this.** These action plans should include commitments to become London Living Wage accredited.
- **Ensure Ealing's career progression pathways are fully inclusive** and that residents from ethnic minorities can access training and qualification routes that will further their career.
- **Examine the EPG in the council's own employees.** Ealing Council publishes its EPG, however this analysis could be bolstered by taking into account additional characteristics such as age, sex, and years of employment, which are all known to be associated with pay, to calculate a 'true' EPG.

Defining the ethnicity pay gap

The central concept behind an EPG is straightforward – **are people from different ethnicities paid different salaries?** That simple question forms the basis for the discussion and analyses presented in this report.

The EPG is clearly analogous to questions about the gender pay gap, which organisations with greater than 250 employees have a statutory obligation to report . At the time of finalising this report[1], there was no statutory requirement on organisations to report EPGs, though the draft Equality (Race and Disability) Bill proposes that ethnicity and disability pay gap reporting will become statutory alongside reporting of the gender pay gap. Against a national figure suggesting that **for every £1 earned by *white* employees, those from ethnic minorities earned 85.2p**[2], many organisations are exploring the publication of ethnicity pay data.

The methodology typically involved in reporting a pay gap is what we will refer to here as a vertical methodology. It identifies the **extent to which different groups are over- or under-represented in different pay bands**[3]. While useful, the vertical approach does not model the complex set of factors known to affect pay such as age, experience, education and qualifications. The risk inherent in vertical analyses of pay gaps is that the variable under consideration (be it gender or ethnicity) might be assumed to play a causal role in the pay gap. However, should this variable be associated with some other variable (or variables) which is the actual cause of the gap, any conclusions drawn about gender or ethnicity are spurious. For instance, if it is the case that education is the driving force behind pay gaps, and ethnic groups differ in their educational attainment, then an observed relationship between ethnicity and pay is spurious.

In the analyses which constitute the body of this report, we adopt what we refer to as a horizontal approach to the ethnicity pay gap (henceforth, the EPG). Broadly speaking, in this approach we **quantify and statistically control for additional variables which may affect the pay gap**. We then examine whether ethnicity still explains variability in pay. The benefits of this are twofold. First, if ethnicity does not explain variability when we control for other variables, then we should focus our efforts on these. Second, if ethnicity does explain pay variability after all other variables have been controlled for, we should consider very different kinds of response.

1. https://www.gov.uk/government/publications/gender-pay-gap-reporting-guidance-for-employers/who-needs-to-report
2. https://www.gov.uk/government/publications/homes-england-gender-and-ethnicity-pay-gap-report-2024/
3. https://www.ealing.gov.uk/downloads/download/4741/gender_pay_gap

The ethnicity pay gap in the UK

The standard quantification of the EPG is to calculate the difference between the median hourly pay of White British and each comparison ethnic group, then divide this difference by the median hourly pay of White British individuals. The result reflects the EPG as a proportion of White British earnings (and leads to the 2024 headline figure of 85.2p). Here, we choose to present actual median pay, noting that expressing differences in terms of percentages often makes them harder to comprehend. The use of absolute median values also avoids ethnocentric reporting.

Pay levels across UK census-defined ethnic groups have changed over time, as illustrated by ONS data spanning the period 2014-2019. Figures 1 and 2 below show the median hourly rate of pay has increased for all ethnic groups over this time period, but disparities remain more or less unchanged. While **earnings in London are higher than in the remainder of England and Wales, the EPG is more marked in the capital than elsewhere.**

Figure 1.
Median pay by ethnicity: England and Wales (excludes London) 2014 - 19. ONS data.

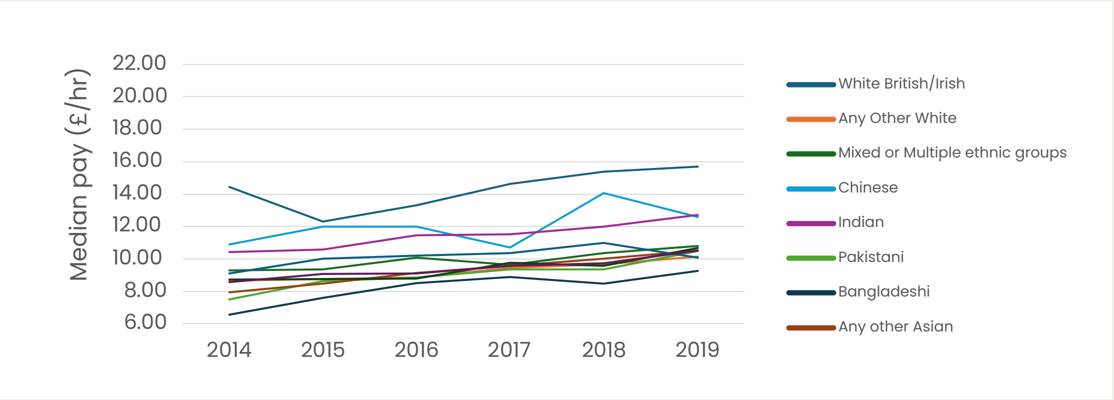

Figure 2.
Median pay by ethnicity: London only 2014 - 19. ONS data

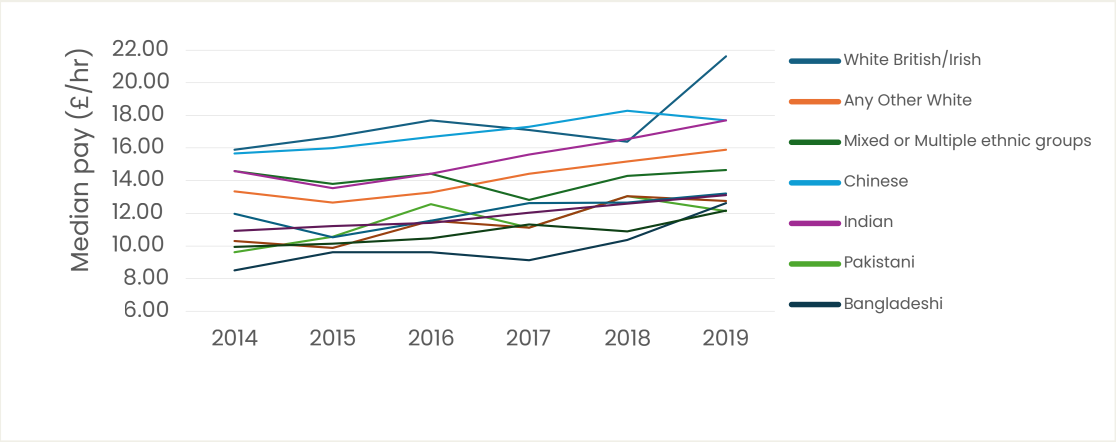

Separate data from the Annual Population Survey (2012-2019) is presented in Figure 3. This survey uses slightly different ethnicity classifications (GSS Harmonised Principle for ethnic groupings) but presents a similar picture. However, due to small sample sizes, data from *White and Black African, Black other, and Arab* ethnicities should be treated with considerable caution as these are unlikely to be representative.

Figure 3. Median pay by ethnicity: England and Wales 2012 - 2019. APS data

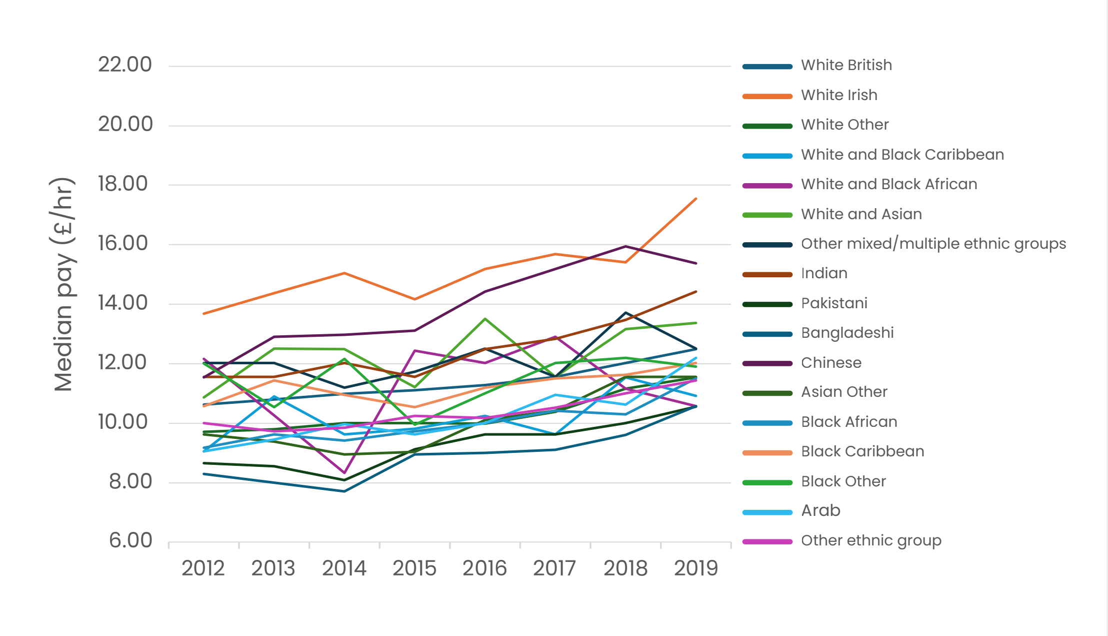

There is therefore good evidence to suggest that **there is an EPG in London, in the rest of England and Wales**, and that this is relatively consistent across different datasets using different classifications of ethnicity.

The ethnicity pay gap across councils in London

Of London's 33 Councils[4] , 24 have published data on their employee EPG either in separate reports or alongside statutory reporting of the gender pay gap. Figure 4 presents mean and median hourly pay gap data. Note that some councils present a negative EPG, meaning that employees from ethnic minorities are more highly paid than the White comparison group. Several councils also report a median EPG of zero with means greater than zero, suggesting skewed data[5]. While only mean and median EPGs are presented here, many councils also include data on additional pay gaps (bonus value, proportion of ethnic groups receiving a bonus, and proportion of staff represented at each salary quartile).

Figure 4. Reported EPG across London Councils.

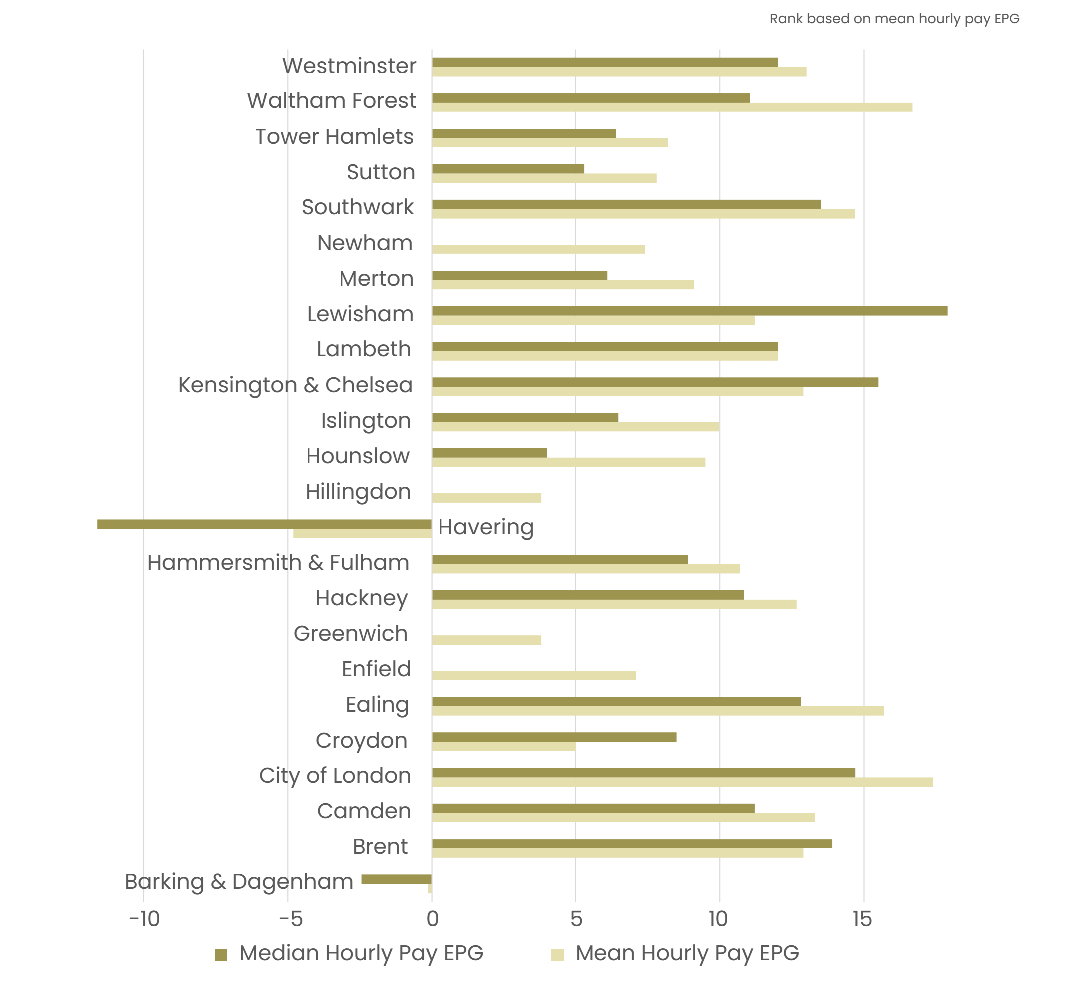

4. Technically 32 plus the City of London.

5. Where medians are less than means the data are positively skewed, meaning the 'hump' in the distribution is towards the lower end, with a long tail at the higher end. This may suggest a small number of individuals in each comparison being paid at relatively high rates.

Across these 24 councils, the EPG is 9.58% for mean hourly pay, and 7.38% for median hourly pay. Using the mean value, this indicates that **ethnic minority employees of councils in London earn 90.42p for every pound their White counterparts earn**. It is important to note the considerable variability in these data.

Causes of the ethnicity pay gap

To address the central question of whether identified EPGs arise solely from ethnicity, more complex statistical analyses than the descriptive data reported above are required. Our principal aim is to identify an outcome (in this case, pay), and statistically predict this using all available data using a technique known as regression. A regression model informs our understanding of which predictors are associated with pay. Knowing these predictors may help us explore opportunities to harmonise pay through targeted interventions. More details and a technical discussion of the forms of regression employed here can be found in Appendix B.

Regression models have been constructed by the Office for National Statistics (ONS) for the APS data presented in Figure 3. At the level of England and Wales, two ethnic groups receive lower pay on average than White individuals. These are *Black, African, Caribbean* or *Black British*: *African, and Mixed or Multiple ethnic groups*: *White and Black African*[6]. However, when country of birth (born in the UK versus born outside the UK) is included as an additional variable, several ethnic groups are paid less than White individuals born in the UK (see Table 1).

Table 1.
Ethnic groups born outside the UK earning less than White individuals born in the UK.

White: Irish
Asian or Asian British: Bangladeshi
Asian or Asian British: Indian
Asian or Asian British: Pakistani
Mixed or Multiple ethnic groups: White and Black Caribbean
Other ethnic group: Arab
Other ethnic group: Any other ethnic group

6. The regressions also reveal expected patterns of results for other variables. Higher levels of pay are associated with living in London, being born in the UK, education to degree-level, a professional occupation, and being male.

This illustrates an important principle. Associations which exist between two variables (for example pay and ethnicity) may arise due to the action of an additional variable or variables. This underpins the importance of examining the most detailed datasets available, as these have the potential to provide critical information in the form of additional variables of interest. Knowing there is an EPG is insufficient if we cannot state with some degree of confidence which other variables (such as educational achievement) are involved[7].

The APS dataset does not permit the kinds of analyses we would ideally like to perform[8]. This is a matter of scale as some of the sample sizes are too low to permit the type of analysis we need. As we add more variables, the combinatorial nature of the dataset expands rapidly. For instance, with just ten categories of ethnicity, five levels of occupation, five of education, two sexes, and five age categories, we create a data structure with 2,500 unique cells. As our analyses require all these cells to be populated with a reasonable number of individuals, only the very largest and most comprehensive datasets permit the kinds of analyses which are required if we are to disentangle the effects of ethnicity from those of other variables.

7. This is similar to but should not be confused with the expression 'correlation does not imply causation'. In these analyses we are only able to make statements about associations (correlations), and not causal links. This weakness is inherent in most work where existing datasets are the primary source of information.

8. Confirmed with the Policy Evidence Analysis team at ONS: Personal Communication, 4th October, 2024.

The London Borough of Ealing

The review above presents a picture of a significant and persistent EPG and identifies some of the key issues in understanding the nature of this. Before turning to an analysis of data specific to the London Borough of Ealing, it is important to discuss what kinds of data exist at the borough level, and how these affect the kinds of analyses which may be performed.

The fundamental unit of analysis linking pay to ethnicity is the individual. Individuals belong to, or identify with, specific ethnic groups, and are paid a certain amount for their labour. Access to individual-level information about pay, ethnicity, and other variables of interest is strictly limited for reasons of confidentiality and sensitivity.

If we cannot access individuals' salaries then we must abstract away from individuals, and group them into higher-order categories. Fortunately, such schemes are well established within ONS datasets, and table 2 outlines the main levels of grouping and how these apply in Ealing.

Table 2. Levels of granularity in data (high to low).

Full title	Short title	Size	Number in Ealing
Individual		1 individual	366,127
Household		1-10 individuals 1 household	137,113
Output Area	OA	100-625 individuals 40-250 households	1,004
Lower Layer Super Output Area	LSOA	1000-3000 individuals 400-1200 households (4-5 OAs)	196
Middle Layer Super Output Area	MSOA	5000-15000 individuals 2000-6000 households (4-5 LSOAs)	39
Towns[9]		Unspecified	7
International Territorial Level 3	ITL3	Unspecified	1

9. This is not an ONS level of organisation, but is important in LBE and is used in some of the analyses which follow.

In general, the greater the number of observations (the sample size), the more powerful the analysis and the more likely we are to detect any differences or associations which exist. Large samples permit us to examine the kinds of complex interactions underpinning the EPG. However, restrictions and limitations on official and public domain reporting mean that some variables are not quantified at the level of the individual or household, and we are therefore forced to use lower levels of granularity if we wish to include these (see Appendix A for a discussion of data limitations). For instance, the Index of Multiple Deprivation (IMD: see Appendix C for a discussion) and its components is available at the level of LSOAs, but not OAs. If we wish to examine the effects of the IMD, we are forced to limit our analysis to the level of LSOAs. This clearly reduces the number of observations in Ealing from 366,127 to 196. Such a reduction in statistical power may render some analyses difficult or misleading to interpret.

In what follows, we adopt an eclectic multi-faceted approach. Analyses are presented at different levels of granularity, with appropriate caveats as to how best to interpret the results. Our aim is to extract meaning from data while adhering to principles of responsible and defensible analysis and reporting. In doing this, we follow guidance on communicating uncertainty issued by the UK Government[10], aiming to present analyses in transparent language, allowing the facts to speak for themselves without overinterpretation, and acknowledging (and ideally quantifying) uncertainty in results. Specifically, each analysis in this report is described with reference to three central questions:

1. **Which data sources were used?**
2. **What do we conclude?**
3. **How confident are we in this conclusion?**

For a technical discussion of how we interpret statistical results please see Appendix E.

10. See https://analysisfunction.civilservice.gov.uk/policy-store/communicating-quality-uncertainty-and-change/

Is there evidence of an ethnicity pay gap in the population of the London Borough of Ealing?

The following section describes the results of analyses based on three main outcomes related to pay. In the first, we examine the components of the Index of Multiple Deprivation (IMD) and ask whether ethnicity is associated with different measures of deprivation. In the second, we examine data on house prices and ask whether these are associated with ethnicity. In the third, we look at low granularity data on income and ask the same question about whether this is associated with ethnicity.

The Index of Multiple Deprivation

Does ethnicity predict income deprivation?

We begin by asking the question of whether income deprivation can be predicted from ethnicity. Is the presence of more or fewer people from particular ethnic groups associated with greater or less income deprivation?[11]

What data are we using?

Census data from the ONS. The variable we aim to predict is income deprivation, at the level of LSOAs, in terms of the number of individuals from each ethnic group who live in each LSOA.

What do we infer?

That income deprivation is greater when there are a) more *Black British African*, b) more *Asian British Indian* people living in an LSOA. The number of *Black British African* people living in an LSOA accounts for approximately 16% of the variation in income deprivation, while the number of Asian British Indian accounts for an additional 9%. No other ethnic groups are associated with increased deprivation.

How confident are we in this inference?

The quality of the data is high, but their specificity is poor. Income deprivation is highly correlated with other measures of deprivation.

11. The analyses reported in this section exclude *White Gypsy* or *Irish Traveller*, and *White Roma* ethnic groups as the average number of individuals across LSOAs in both groups is <20.

Does ethnicity predict employment deprivation?

An observation from the LBE data is that there is a very high association between income and employment deprivation. This association is so high that it is effectively impossible to disentangle the effects of these aspects of deprivation. We therefore repeat the above analysis using employment deprivation as the outcome.

What data are we using?

Census data from the ONS. The variable we aim to predict is employment deprivation, at the level of LSOAs, in terms of the number of individuals from each ethnic group who live in each LSOA.

What do we infer?

That employment deprivation is greater when there are more *Black British African people* living in an LSOA. **The number of *Black British African people* accounts for 12% of the variation in employment deprivation.** No other ethnic groups are associated with increased deprivation.

How confident are we in this inference?

The quality of the data is high, but their specificity is poor. As employment and income deprivation are so highly related, it is not possible to disentangle their effects in this dataset.

Does ethnicity predict education deprivation?

A key third factor (and one which also correlates highly with income deprivation) is education deprivation. We therefore repeat the above analysis with this as the outcome.

What data are we using?

Census data from the ONS. The variable we aim to predict is education deprivation, at the level of LSOAs, in terms of the number of individuals from each ethnic group who live in each LSOA.

What do we infer?

That education deprivation is greater when there are more *Black British African* people living in an LSOA. The number of *Black British African* people accounts for 12% of the variation in education deprivation. No other ethnic groups are associated with increased deprivation.

How confident are we in this inference?

The quality of the data is high, but their specificity is poor.

The three analyses in this section paint a very similar picture – that the number of *Black British African* people living in an LSOA predicts levels of income, employment, and education deprivation. Around 12-16% of the deprivation across these three measures is accounted for by an ethnic group which constitutes approximately 6% of Ealing's total population. Additionally, income deprivation is predicted by the number of *Asian British Indian* people in each LSOA.

One important caveat of these analyses is that the IMD is a measure of deprivation, and consequently inferences drawn from these outcomes should be interpreted with caution. The IMD is not a measure of affluence, as the relative amount of deprivation within an LSOA says nothing about its relative affluence. An absence of high levels of deprivation is not evidence for the presence of high levels of income, employment, and education. As such, it is important to augment these analyses with other measures of affluence.

Deprivation across Ealing's seven towns

While LSOAs are a creation of the ONS and represent a useful and consistent way to identify and quantify similar-sized neighbourhoods across the UK, the seven towns in Ealing (Acton, Ealing, Greenford, Hanwell, Northolt, Perivale, and Southall) are historical, indeed ancient settlements[12]. As such, they represent an important local level of organisation which merits investigation.

In the following section we seek to answer the question of whether, and to what extent, levels of deprivation differ across the seven towns. Rather than predicting relationships, we turn our attention to detecting differences between the towns. Please refer to Appendix D for a technical discussion of the analyses reported here.

We begin by examining the towns' ethnic makeup, move to identifying differences in deprivation levels, and finally discuss the more complex question of how these differences may be associated with ethnicity.

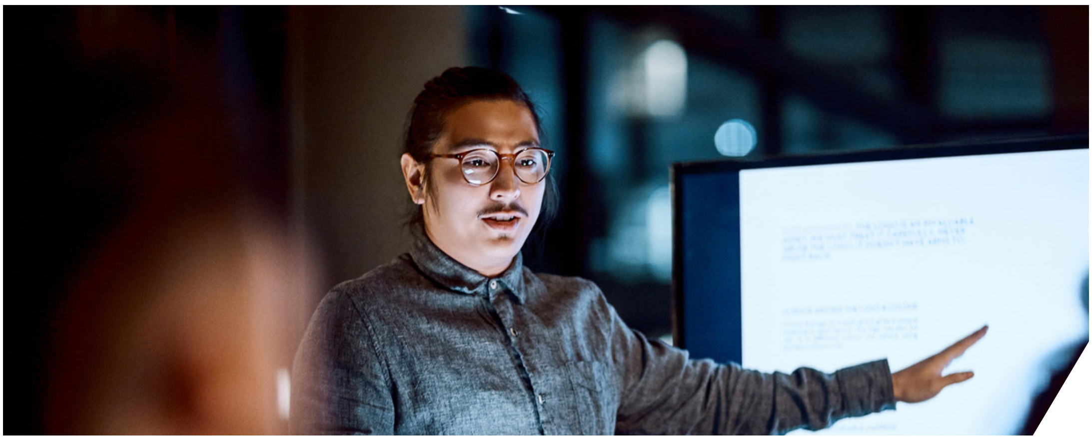

12. We are indebted to the GIS team at LBE for their assistance in mapping LSOAs onto the 7 towns.

Ethnicity populations in the seven towns

Figure 4 presents population data across the seven towns broken down by ethnicity. The mapping of LSOA to towns is not a direct one to one mapping as LSOAs are not created with reference to town boundaries. Figure 5 illustrates this issue where an LSOA identified as being in Southall crosses the boundary with Greenford. In defence of these analyses, we note that the data on which Figure 4 is created map well onto published data for town populations, suggesting the error involved in this organisation is small.

Figure 4. Town population by ethnicity

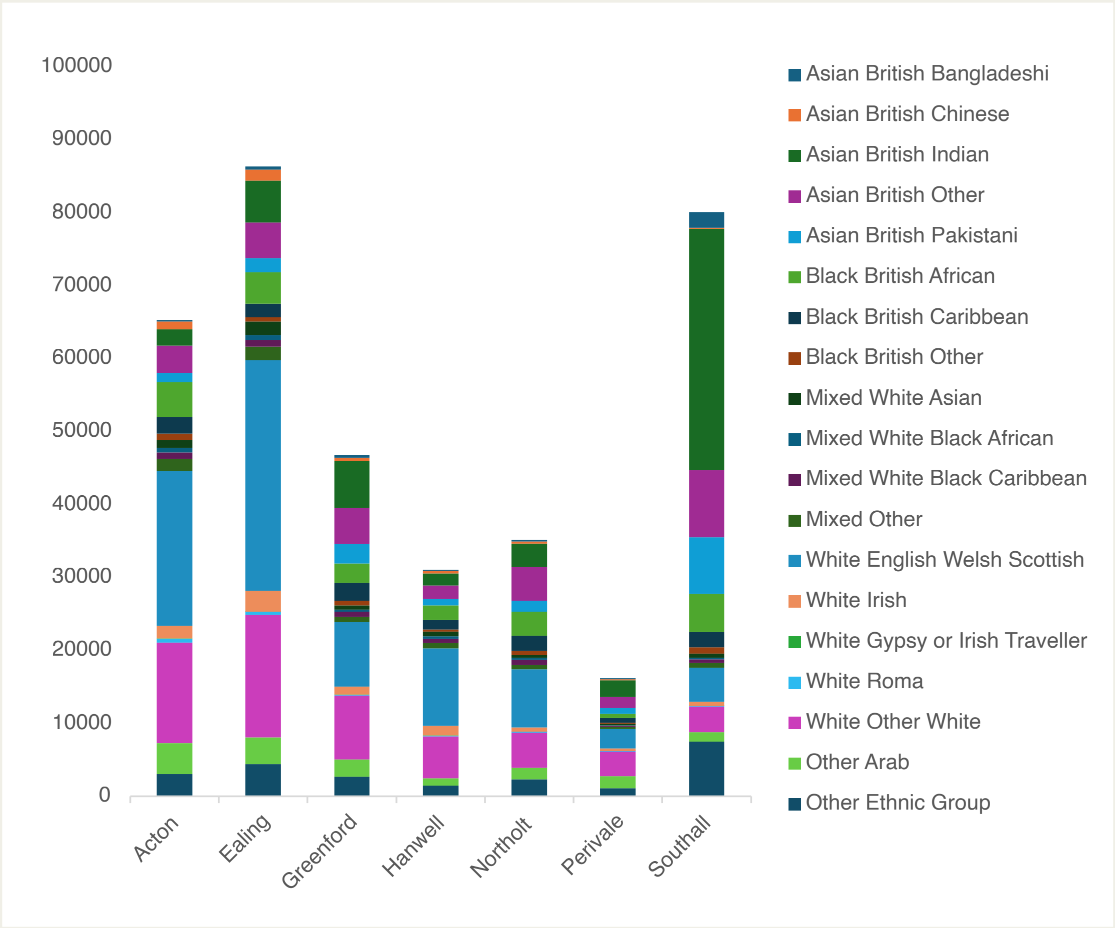

Figure 5. Example discrepancy between LSOA and Town boundaries

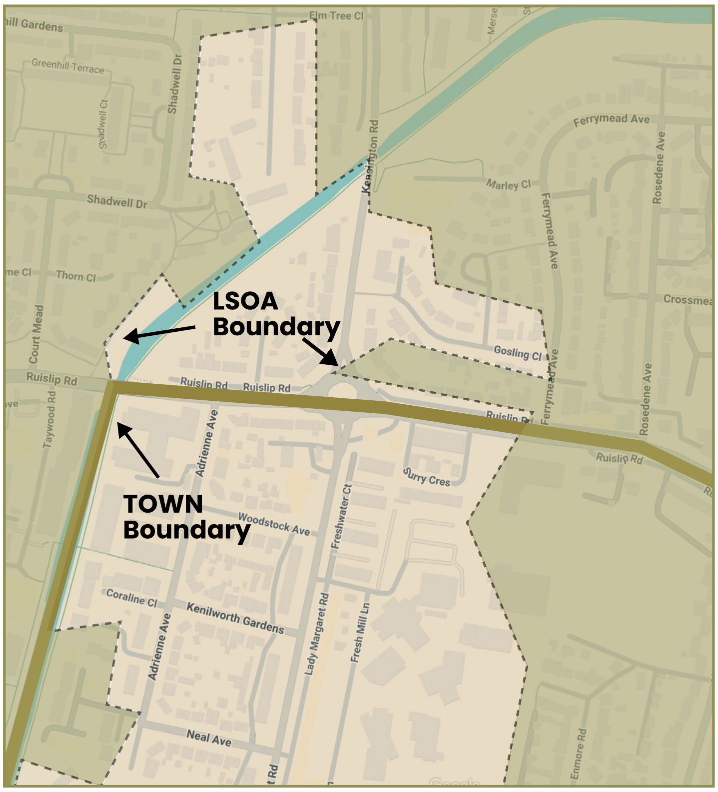

Does income deprivation differ across towns?

Figure 6 illustrates the average levels of income deprivation in the seven towns (higher levels indicate higher levels of deprivation[13]).

Figure 6. Income deprivation

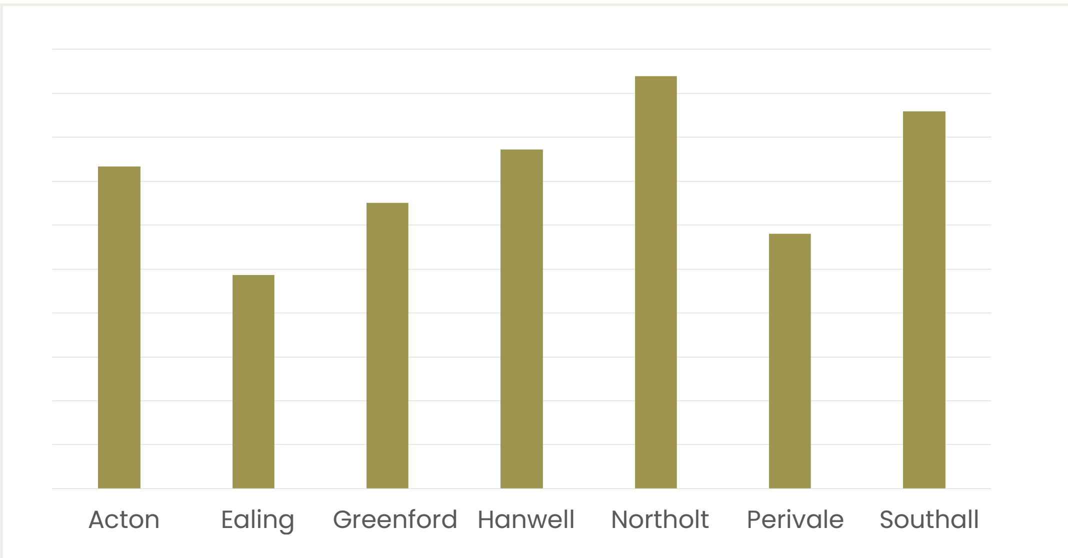

13. We have also removed values from the y axis as deprivation scores cannot be directly compared across domains of income, employment and education. The analyses presented here focus on relative deprivation.

What data are we using?

Census data from the ONS, augmented by a mapping of LSOAs onto the seven towns. We examine whether there are differences in income deprivation across the seven towns, and which towns differ from one another.

What do we infer?

There are large differences in the levels of income deprivation. There is a 'cluster' of six towns (Greenford, Hanwell, Acton, Northolt, Perivale and Southall) which have statistically similar levels of income deprivation ranging from 12-19%. **Ealing has lower income deprivation levels than Acton, Hanwell, Northolt, and Southall.**

How confident are we in this inference?

The quality of the data is good, though there is an imperfect match between LSOAs and the seven towns. The specificity of the data is poor. The variability contained in the data is high.

Does employment deprivation differ across towns?

Figure 7 illustrates the average levels of income deprivation in the seven towns (higher levels indicate higher levels of deprivation).

Figure 7. Employment deprivation

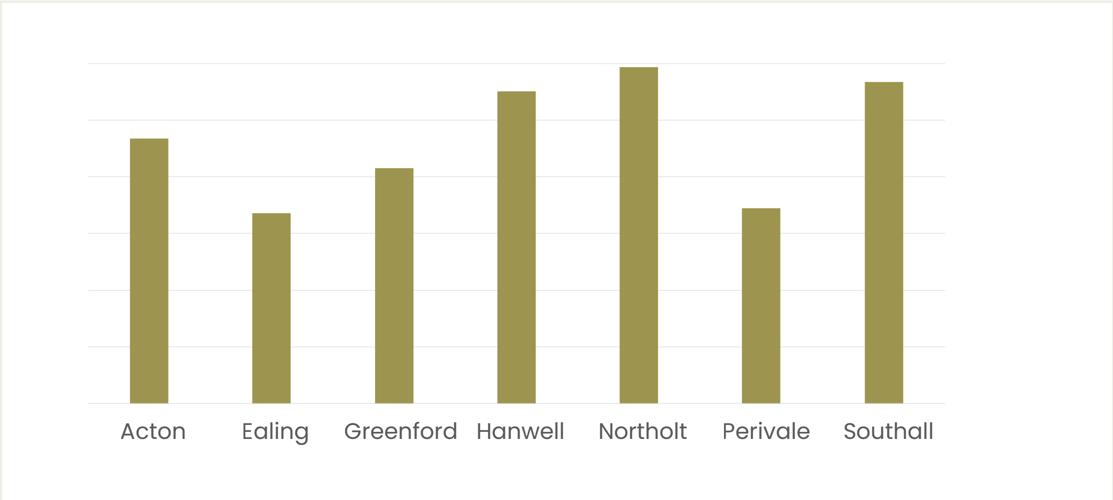

What data are we using?

Census data from the ONS, augmented by a mapping of LSOAs onto the seven towns. We examine whether there are differences in employment deprivation across the seven towns, and which towns differ from one another.

What do we infer?

There are large differences in employment deprivation. The same cluster of six towns emerges with employment deprivation scores ranging from 8-12%. **Ealing has lower employment deprivation than Hanwell, Northolt, and Southall.**

How confident are we in this inference?

The quality of the data is good, though there is an imperfect match between LSOAs and the seven towns. The specificity of the data is poor. The variability contained in the data is high.

Does education deprivation differ across towns?

Figure 8 illustrates the average levels of income deprivation in the seven towns (higher levels indicate higher levels of deprivation).

Figure 8. Education deprivation

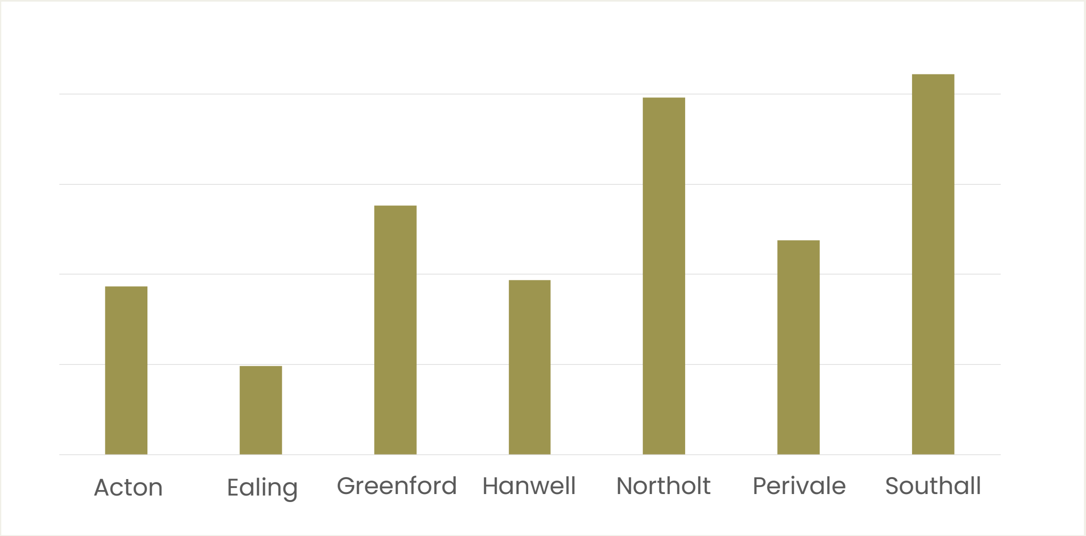

What data are we using?

Census data from the ONS, augmented by a mapping of LSOAs onto the seven towns. We examine whether there are differences in education deprivation across the seven towns, and which towns differ from one another.

What do we infer?

There are large differences in the levels of education deprivation. **Ealing has less education deprivation than Acton, Greenford, Northolt, Perivale and Southall**. Acton, Hanwell, Greenford, and Perivale are indistinguishable. **Southall and Northolt are more educationally deprived than all other towns** and are indistinguishable from each other.

How confident are we in this inference?

The quality of the data is good, though there is an imperfect match between LSOAs and the seven towns. The specificity of the data is poor. The variability contained in the data is high.

It is clear from the above data and analyses that there are large differences in the ethnic makeup of the towns' populations, and that income, employment and education deprivation vary significantly. The degree to which different ethnic populations are associated with different kinds of deprivation across the seven towns is a straightforward question, but one which presents serious challenges.

With 196 LSOAs in the borough, the number of LSOAs identified as 'belonging' to each town is 28 (there is variation as the population, and by implication the number of LSOAs, varies across towns). The granularity of the resulting dataset is therefore small (~28) while the number of predictors remains high (19 ethnic groups). The results of analyses where the number of predictors is high relative to the number of cases are inherently misleading, and we do not therefore predict deprivation within towns as a function of high-granularity ethnicity.

Nevertheless, we believe the interplay between ethnicity and deprivation within each town represents potentially useful information and therefore collapse ethnicity into two categories and adopt the distinction between White and ethnic minorities in what follows.

Does ethnicity predict income deprivation in the seven towns?

What data are we using?

Census data from the ONS, augmented by a mapping of LSOAs onto the seven towns, incorporating a binary classification of ethnicity as White or ethnic minorities. We examine whether ethnic population predicts income deprivation.

What do we infer?

In Ealing, Northolt, Perivale, and Southall there is no relationship between income deprivation and the size of either the *White* or ethnic minorities populations.

In Greenford and Hanwell, a higher ethnic minority population is associated with higher income deprivation. There is no association between income deprivation and the White population.

In Acton, a higher White population is associated with lower income deprivation, and **a greater ethnic minority population is associated with higher income deprivation.**

All of the reported effects are small.

How confident are we in this inference?

The quality of the data is good, though there is an imperfect mapping between LSOAs and the seven towns. The specificity of the data is poor, and the number of outcomes is low. The classification of ethnicity into two categories is driven by data considerations and its granularity is extremely low.

House prices

A second proxy of pay is housing prices. Again, this is an approximate proxy, but it has the advantage of being available at the level of LSOAs. This therefore permits us to rigorously investigate whether the number of individuals from each ethnic group is associated with house price.

Does ethnicity predict house prices?

What data are we using?

Census data from the ONS identifying median house price in each LSOA in June 2019. We examine whether house price can be predicted by the number of people from each ethnic group.

What do we infer?

Ethnicity predicts a large proportion of variation in house prices. The only ethnic group which statistically predicts house prices is *White English Welsh Scottish*. **Higher house prices are associated with a higher number of White English Welsh Scottish living in the LSOA.**

How confident are we in this inference?

The quality of the data is good, though 20 LSOAs are not included in the analysis as there is no available house price reported for June 2019. The specificity of the data is poor.

Household income

Data on income are publicly available at the level of MSOAs (areas of 5000-15000 people) and are therefore of very low granularity. Due the relatively small number of MSOAs in LBE (39), and some very high associations between the numbers of different ethnicities living in MSOAs, the regression approach used in the previous section cannot be applied. Instead, we examine standard correlations between the number of people from each ethnic group, and the overall household income in the MSOA.

Does ethnicity predict household income?

What data are we using?

Census data from the ONS. We examine the association between household income at the level of MSOAs and the number of individuals from each ethnic group who live in each MSOA.

What do we infer?

Three ethnic populations are associated with higher household income. That is, the greater the number of people from these groups who live in an MSOA, the greater the average household income is in that MSOA.

Mixed White Asian
White English Welsh Scottish
White Irish

Five ethnic populations are not associated with household income. That is, there is no link between the number of people from these groups who live in an MSOA, and the average household income in the MSOA.

Asian British Chinese
Mixed White Black African
Mixed Other
White Other
Other Arab

Eight ethnic populations are associated with lower household income. That is, the greater the number of people from these groups who live in an MSOA, the lower the average household income is in that MSOA.

Asian British Bangladeshi
Asian British Indian
Asian British Pakistani
Asian British Other
Black British African
Black British Caribbean
Black British Other
Mixed White Black Caribbean

How confident are we in this inference?

The quality of the data is high, but their specificity is poor and the granularity is low. The analysis ignores all relationships between the numbers of individuals from different ethnic groups, as it is simply a series of correlations. Notwithstanding this, the effects are large, and there is a relatively clear distinction between those which show a relationship with household income, and those which do not.

Conclusions and recommendations

This project had two main aims. First, to examine the nature of the ethnicity pay gap in the population of the London Borough of Ealing, and second to go beyond this and consider how ethnicity and other factors contribute to this. This final section summarises our findings and presents recommendations for further work.

The EPG is present in one form or another in every dataset we examined. We see evidence of certain ethnic populations within local neighbourhoods being associated with greater deprivation in terms of income, employment and education. In particular, ***Black British African*** **people, who constitute a population of approximately 23,000 across the borough appear to experience greater deprivation than other groups.**

At more general levels of geography, large areas of the borough show patterns where **the greater the White population, the greater the average household income is**, whereas the populations of many other ethnic groups predict lower levels of household income.

At a more general level still, there are important differences across the seven towns which constitute the borough, and some evidence (though this requires a significant caveat as the measure of ethnicity we were able to use was very crude) that **in Acton, Greenford, and Hanwell, higher income deprivation is associated with larger ethnic minority populations.**

The limitations of the various datasets utilised here make it impossible to draw causal conclusions. We do not know whether the EPG is due to ethnicity, or some other set of factors, but it remains clear that across a number of different datasets, different measures, and different levels of analysis, some ethnic groups fare less well than others.

Throughout this discussion, the epistemological 'elephant in the room' has been the lack of availability of data which are direct measures of pay. While this is a feature of public datasets, it presents an important limitation to all that has been discussed. When pay is equated with affluence, we should be careful to draw strong conclusions from deprivation data, though we may express greater confidence when we conceptualise pay as relating to the absence of deprivation.

The new Equality (Race and Disability) Bill will make ethnicity pay reporting compulsory for employers with at least 250 employees and improve access to data to further bolster the findings of this report.

The recently published London Growth Plan has a prime objective to reduce inequality for Londoners, with a new inclusive talent strategy to set out how marginalised groups can be better supported into well-paid employment. This report will be shared with the Greater London Authority as evidence to support the production of that strategy.

Recommendations

- **In preparation for the Bill, ensure large employers are ready to meet this statutory requirement and encourage them to publish reports in advance.** This could require targeted engagement, guidance and specialist business support through the council's Good for Ealing programme.
- **Encourage employers of less than 250 employees to analyse and report on their EPG voluntarily.** This is particularly relevant in Ealing, where micro, small, and medium businesses make up 99.8% of the local economy. This could require a communication campaign, guidance and specialist business support through the council's Good for Ealing programme.
- **Support businesses of all sizes to develop actions plans to address their EPG**. This could require guidance and specialist business support through the council's Good for Ealing programme.
- **Continue to work with Ealing's London Living Wage** group to tackle pay inequality across the borough, including targeted efforts to increase the number of London Living Wage accredited businesses in parts of the borough where there are more Black British African or Asian British Indian people living and income deprivation is greater.
- **Through Ealing's new growth sector forums,** explore how the EPG varies across the foundational and growth sectors that makeup Ealing's economy, and measures that need to be put in place to address this.
- **Through the council's purchasing powers,** encourage more businesses to report on and develop action plans to address the EPG. This should be a key consideration of the social value component of a procurement.

- **Encourage all commercial tenants of council property to report their EPG and develop action plans to address this.** These action plans should include commitments to become London Living Wage accredited, demonstrating their commitment to fair pay.

- **Ensure Ealing's career progression pathways are fully inclusive** and that residents from ethnic minorities can access the training and qualification routes that will further their career.

- **Examine the EPG in the council's own employees.** Ealing Council publishes its EPG, which is slightly higher than the average (both mean and median hourly pay) of the other 23 boroughs for whom data are available. The vertical mean and median hourly pay rates (and the associated proportion of White and ethnic minority employees represented in the four ascending pay bands) are important and useful information. However, these could be significantly augmented through the use of a horizontal analysis along the lines of the approach adopted in this report. By taking into account additional characteristics such as age, sex, and years of employment which are all known to be associated with pay, it should be possible to calculate a 'true' EPG. People are paid differently for the jobs they do, for a variety of reasons. Experience and performance might be very good reasons for this, while sex and ethnicity are not. Teasing apart the relative contributions of all the factors affecting pay is an important and worthwhile challenge.

Appendices

Appendix A: data sources and challenges

The EPG can be examined at a number of levels, including national, regional, local, and organisational. Each level is associated with datasets containing different measures conceptualised in specific and sometimes idiosyncratic ways. Some data (for example the ethnic composition of the United Kingdom) is public domain, whereas others (e.g. salaries of individual employees within an organisation) are restricted. Both pay and ethnicity are measured and defined in a variety of ways, and we examine the challenges this presents to understanding the EPG below.

Measuring pay

Perhaps the best, and consequently most widely used proxy for pay is self-reported income (both individual and household). While such measures are relatively accurate and reliable, they often lack precision, as people are presented with a limited number of options (for example, £0-£20,000, £21,000-£40,000, and so on). This has consequences for the kinds of analyses which can be performed on the data, in terms of both its nature and precision.

In the absence of direct or self-reported income data, a number of additional proxies can be identified, each with limitations of accuracy and applicability. These include such measures as average house values and estimated household income. A major challenge with using such measures as proxies is that they may in fact be responsible for any identified gap. For instance, the known association between educational level and pay might lead us to substitute educational level as a proxy for pay. In doing so, we lose the opportunity to determine whether education level, rather than ethnicity, is determining the EPG, and miss its potential significance (in this case, of targeting enhanced educational opportunities).

Categorising ethnicity

As an evolving and socially important concept, ethnicity has been captured in many different ways. Even within datasets, categories change over time because of increased understanding of the complexity of ethnicity and changes in the makeup of the population in question. Ethnicity may also be replaced with nationality.

Broad categories of ethnicity are simple to apply, and can provide useful information, but more fine-grained classifications are critical. For instance, there is a general move away from using the label BAME (Black And Minority Ethnicity) as this is highly heterogenous.

In Ealing, this is of significance for several reasons. Ealing is ethnically diverse, and contains the largest Polish, Afghan, and Serbian populations in England and Wales, the second largest Japanese and Iranian populations, and the fourth highest Arab population. To the extent that these populations (and all smaller populations represented within the borough) present different patterns of employment, it is important to gather intelligence at the level of nationality rather than some more abstract ethnic grouping. However, datasets generally only encode ethnicity.

In the data and analysis which is used in this report, we strive to present the most fine-grained analysis available. For instance, if a dataset includes 5 and 18 categories of ethnicity, we aim to present the latter. We also adopt the convention of italicising ethnic groups, so *Other ethnic group: Arab* represents the official category adopted by the relevant dataset.

Appendix B: technical discussion of regression models

Regression is a complex statistical technique, having many variants. In all, the aim is to predict an outcome(s) based on predictors. The analysis combines the predictors in a way which best captures the relationship, and determines which predictors play a role in predicting the outcome.

When datasets consist of outcome and predictor variables that can be combined mathematically, the form of regression known as linear regression is preferred. In this technique, the numbers representing the quantities have a real value, and can be combined together, multiplied, subtracted, etc. For instance, variables such as 'pay', when measured as a single point value, are 'proper' numbers, and can be used in linear regression. If, however, we measure 'pay' using a series of categories, as is common in census work, we lose some of its mathematical characteristics. While we know that a person who earns 'between £21,000 and £40,000' earns less than someone who earns 'between £41,000 and £60,000', we don't know exactly how much less they earn. As linear regression makes use of these calculated differences, we cannot use it when the data do not support this kind of operation.

When datasets contain a wide variety of variables, forms of regression called logistic regression, or ordinal regression are typically employed (the main difference between them being whether the outcome we are trying to predict is a binary outcome, such as sex, in which case we use logistic regression, or is ordinal, such as salary scales, in which case we use ordinal regression). The aim of these forms of regression is the same – how to predict an outcome. Unlike linear regression, they are not dependent on all variables being quantities and can deal with categorical data (e.g. sex), and ordinal data (e.g. salary scales from census data). One key aspect of logistic and ordinal regression is that all results are interpreted relative to a reference category. For instance, setting the reference category for ethnicity to 'White British' arbitrarily allocates a value of 1 to the outcome of White British individuals in the dataset. The regression then computes the outcomes of all other ethnicities relative to this. An ethnic group with a higher level on the outcome than White British individuals would have a score greater than 1, and a group which scored less on average than White British individuals would score less than 1.

In selecting a regression method to use, the principle is to always use the most powerful legitimate method. Linear regression is more powerful than logistic or ordinal regression, and we therefore choose to use this whenever possible. Should our variables not correspond to the requirements of linear regression, we choose either logistic or ordinal regression based on the nature of the outcome variable.

Appendix C: The Index of Multiple Deprivation

The IMD is managed by the Ministry of Housing, Communities and Local Government, and consists of seven components, combined to generate a single score for each Lower-layer Super Output Area in England. The seven components, and their relative weightings, are:

- Income (22.5%)
- Employment (22.5%)
- Education (13.5%)
- Health (13.5%)
- Crime (9.3%)
- Barriers to housing and services (9.3%)
- Living environment (9.3%)

Data for the IMD is available in terms of the absolute score for each component (derived from a number of different variables totalling 39 across the seven components), rank within England (out of 32,844, where lower ranks equate to higher levels of deprivation), and deciles (where the components and overall IMD are split into 10 equally sized groups, with 1 being the most deprived).

Official recommendations are that ranked data should be used due to their greater level of precision. The absolute scores for each LSOA are provided to three decimal places (rounded, and therefore consisting of 1000 discrete values), and are therefore less precise than the ranked data, which include 32,844 discrete values.

However, we note that absolute scores permit direct numerical calculations, while ranked scores are ordinal. Although the ranked scores are indeed more precise, we cannot make statements about the magnitude of the difference between any two ranks. We therefore use absolute scores in our analyses, as this supports the use of more powerful statistical tests.

Appendix D: technical discussion of difference models

We are often interested in detecting whether certain identified groups differ in terms of some outcome. When the nature of the data permit, the general statistical procedure known as analysis of variance (ANOVA) is an appropriate method to use.

In essence, the underpinning assumptions and calculations of ANOVA are identical to those of regression (they are both variants of the general linear model which is a set of assumptions about the nature of data). While regression focusses on prediction, ANOVA focusses on differences. As with many statistical techniques, ANOVA comes in a number of different forms and provides relatively standard metrics which help us make up our minds about whether there really are differences between our groups, and how large these are.

ANOVA is frequently referred to as an 'omnibus' test. It can compare any number of different groups with one another and determine whether they differ in some way. What it does not tell us is which specific differences exist. For instance, if we compare income deprivation across the seven towns in the borough and find an effect, we know that some of the towns differ in their income deprivation, but not which ones.

A technique known as post hoc testing allows us to tease apart these additional effects. Briefly, the technique compares each possible pairing and determines whether the two members of the pair differ or not. As this can involve a potentially large number of comparisons (in the case of the seven towns, there are 21 different comparisons), results need to be statistically adjusted. Here, we choose to be statistically conservative – that is, we adopt a particular form of adjustment (known as a Bonferroni correction) which reduces the chances of drawing erroneous conclusions about differences when they do not in fact exist. We believe this is defensible when data considerations include ambiguity and a lack of measurement precision.

The end result of this process is that we can state with a quantified level of confidence whether two groups differ on some measure, estimate how large this difference is, and describe the precision of this estimate.

Appendix E: technical discussion of statistical findings

Statistical tests typically provide three separate pieces of information which collectively inform our thinking. These are measures of statistical significance, effect size, and confidence intervals, and each provides critical information. Each is discussed below.

Statistical significance

Statistical significance is the most commonly reported outcome of statistical tests. It is also the least useful, most controversial, and most poorly understood. When a test is reported as 'significant', usually with an associated probability value (p) which is less than some specified value (usually .05, or 5%), the effect being measured by the test is treated as a 'real' effect. That is, something which is actually happening, and not merely an accident of the data. As numerous statisticians have pointed out, this is an error.

A p value is in fact a special kind of probability called a conditional probability. Conditional probabilities are probabilities of one thing being likely given that another thing is true or assumed to be true. In the case of statistics, the thing we assume to be true is typically referred to as a null hypothesis. The null hypothesis generally states that whatever difference or association we are interested in does not actually exist, and that we are wrong.

Why is this so useful? The utility of a null hypothesis is that we know its properties (e.g. that two groups are the same on the measure of interest). Furthermore, due to the underlying properties of data, we also know that when we collect data, these collections (or samples) have specific properties. One important property of a sample is that we can calculate how likely or unlikely it is, given that the null hypothesis is true.

As an example, let's say we are interested in examining whether there are differences in income deprivation across the seven towns in the borough of Ealing. The null hypothesis is straightforward to articulate – there are no differences in income deprivation across the seven towns in Ealing. We can then calculate the levels of income deprivation in each of the seven towns, compare these, and produce a p value which tells us how likely our results are on the assumption that there are in fact no differences in income deprivation.

The sleight of hand which follows is tricky, but critical. If the probability of obtaining the data we have (i.e. actual measured levels of income deprivation across the seven towns in Ealing) is very low (say, less than 5%), we conclude that in fact there are differences in income deprivation across the seven towns in Ealing.

While statistical significance is very widely used, its conclusions are often erroneous, and it is crucial to augment it with other information.

Effect size

When we talk about 'significance' in statistics we really mean confidence – do we believe that something other than randomness is actually happening in our data. If we satisfy ourselves that something is indeed going on, then it is important to understand the magnitude of whatever it is we are measuring. We refer to this as an effect size.

An effect size is conceptually very simple – it is a measure of how large the effect is. There are several ways to quantify effect sizes, but all are normally transposed onto a simple description of effects as small, medium, or large. By knowing how 'big' an effect is, we can make sensible decisions about priorities, whether to invest resources in changing an effect, and so on. These decisions are informed by the effect size, but are also guided by policy, politics, economics, and a host of external forces.

Confidence interval

Our final piece of the statistical puzzle is another unfortunate mislabelling. A confidence interval is really a statement about uncertainty. Formally, a confidence interval states that if you repeated an analysis many times, and calculated a confidence interval for each sample, the true difference would be contained in 95% of these intervals.

An easier (though technically incorrect) way to think about a confidence interval is that it provides a range of values within which the 'true' effect probably lies. Hence, it is a statement about uncertainty.

One important and useful rule is that if the confidence interval includes 0 (i.e. the range of possible outcomes is a negative number at one end of the scale, and a positive number at the other), this suggests that the 'true' effect is in fact no effect. A confidence interval containing a 0 should properly be considered as no effect at all.

Our approach

In discussing the findings of statistical tests, we adopt the following conventions.

i) We consider an effect to be 'real' when it is associated with a significant statistical result and the confidence interval does not cross zero.

ii) Where appropriate we describe the size of the effect in terms of the three simple labels, small, medium, or large. When it is important to do so, we will expand on what the implications of these may be.

iii) We include a brief statement about the quality and specificity of the data used in the analysis.

UNIVERSITY OF WEST LONDON

Ealing

www.ealing.gov.uk

CREATING GOOD JOBS

TACKLING THE CLIMATE CRISIS

FIGHTING INEQUALITY

Good for Ealing