The Scandinavian
Journal of Economics

Scand. J. of Economics 125(1), 73–106, 2023
DOI: 10.1111/sjoe.12511

Where is the Land of Hope and Glory? The geography of intergenerational mobility in England and Wales*

Brian Bell†

King's Business School, London WC2R 2LS, UK
brian.bell@kcl.ac.uk

Jack Blundell‡

Stanford University, Stanford, CA 94305, USA

Stephen Machin§

London School of Economics and Political Science, London WC2A 2AE, UK

Abstract

We present a new analysis of intergenerational mobility across three cohorts in England and Wales using linked decennial census microdata, focusing on occupation, homeownership, and education. Four main results emerge. First, area-level differences in upward occupational mobility are highly persistent over time. Second, measures of absolute and relative mobility tend to be spatially positively correlated. Third, there is a robust relationship between upward educational and upward occupational mobility. Last, there is a small negative relationship between upward homeownership mobility and upward occupational mobility, revealing that social mobility comparisons based on different outcomes can have different trends.

Keywords: Intergenerational mobility

JEL classification: *I*24; *J*62; *R*23; *R*31

---

*We thank three anonymous referees for very helpful comments. This research was partly funded by the ESRC at the Centre for Economic Performance (CEP) at the London School of Economics and Political Science (LSE). J. Blundell is grateful for the generous support of the Europe Center at Stanford University. Permission from the Office for National Statistics (ONS) to use the Longitudinal Study is gratefully acknowledged, as is the help provided by staff of the Centre for Longitudinal Study Informational and User Support (CeLSIUS). CeLSIUS is supported by the ESRC Census of Population Programme (Award Ref: ES/R00823X/1). The authors alone are responsible for the interpretation of the data. This work contains statistical data from ONS that is Crown Copyright. 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 datasets that may not exactly reproduce ONS aggregates.

†Research Associate at the CEP.
‡Visiting Fellow at the CEP.
§Director of the CEP.

1. Introduction

There is copious evidence showing that life chances are affected by family background.[1] Children whose parents went to university are themselves more likely to do so, and they are more likely to be higher earners if their parents were; and it is the same story for homeownership. What is far less clear, and for which there is almost no evidence for England and Wales, is whether the role of family background is attenuated or enhanced by the place where children grow up. In other words, is the link between child and parental university attendance the same in Liverpool as in Lambeth? If your father was a low-skilled worker, does the probability that you become a high-skilled professional depend on whether you grew up in Oldham or Oxford? Such questions are both intrinsically interesting for academic research and at the same time crucial for policy. England and Wales offer a particularly interesting setting for looking at spatial mobility patterns and their evolution through time. There has been immense academic and policy discussion surrounding intergenerational mobility, with much debate over aggregate patterns. The countries also exhibit significant regional imbalances. Akin to the US, there is a dominant narrative of regions being "left behind", as other areas excel. For example, Inner West London is the richest area in Northern Europe, whereas West Wales is the poorest (Eurostat, 2019).

Much of the existing evidence on intergenerational mobility in the UK focuses on relatively small samples of individuals in longitudinal studies that follow a group of people born in the same week in a particular year (Blanden et al., 2004; Blanden and Machin, 2008). These cohort studies are infrequent, suffer from significant attrition over time, and do not provide a large enough sample to provide estimates at anything other than the national level. The UK's Social Mobility Commission (SMC) produces a social mobility index for areas each year; however, the major components of this index focus on wages, employment, schooling, occupation, and homeownership for all adults living in an area, irrespective of background. As acknowledged in Social Mobility Commission (2017), this is a limited proxy of spatial differences in intergenerational mobility.

There are two recent survey-based papers investigating regional differences in intergenerational mobility in the UK. Friedman and Macmillan (2017) study Understanding Society and the Labour Force Survey to estimate rates of mobility across 19 areas of the UK. Exploiting the large sample size of these surveys, they find substantial differences in mobility across regions. While a valuable contribution, an important data shortcoming of this work

[1] For reviews of the economic literature on intergenerational mobility, see Black and Devereux (2011). The vast sociology literature on the topic is discussed in Bukodi and Goldthorpe (2018) and a more populist discussion is given in Elliot Major and Machin (2018).

is that it has to look at individuals' location as adults, rather than where they grew up. This makes it less useful for understanding the geographical origins of differences in rates of mobility. It also uses retrospective reports of parental status, so it will potentially suffer from recall error. Similarly, Rohenkohl (2019) uses longitudinal household survey data to estimate regional intergenerational mobility for 12 regions of the UK. The full length of the panel is exploited so that recall error is not an issue, and the author is able to explore pre- and post-transfer household income rather than just individual labour market earnings. The paper demonstrates a strong North−South divide in mobility rates. However, the sample sizes available for regional estimates are limited, forcing a focus on highly aggregated regions. Neither of these studies is able to investigate changes over time.

In this paper, we significantly improve on both of these dimensions by using linked decennial census data contained in the Longitudinal Study of England and Wales (LS) to study much more disaggregated geographies and trends of intergenerational mobility.[2] The LS data are a 1 percent sample from the decennial census and life events data. They can be structured so as to follow three cohorts of individuals born in 1954−1963, 1964−1973, and 1974−1983 over time. Parental occupation, education, and homeownership are observed when the individual cohort members were aged 8−17 and living with their parents. By first observing the cohort members when they are under 18 and living at home allows us to directly link parent and child. The same outcomes can be looked at for the cohort members when aged 28−37 in the 1991, 2001, and 2011 censuses.

A key feature of this paper is that these data offer scope to focus on multiple measures or dimensions of mobility; see the recent review of Deutscher and Mazumder (2021) for an extensive discussion of a large number of metrics of absolute and relative mobility. The LS data make this possible, and we argue that this is a useful approach for two reasons. First, we might consider that socio-economic status cannot be reduced to a single measure such as income. Rather, there exists some latent socio-economic status that is noisily measured by multiple observables, such as income, wealth, health, or education. Alternatively, one might argue that there does not exist some latent status variable, but that there are simply multiple dimensions of welfare, which cannot be combined without making strong assumptions. We consider both these justifications to be valid, and present separate results for each dimension of mobility. One could combine them into an index, but we do not

[2] As the census is administered separately for Scotland and Northern Ireland, these areas are excluded from this dataset. As of 2015, these omitted countries constitute approximately 11 percent of the UK population. For brevity, we refer to combined estimates for England and Wales as "national" results.

do that here, in part – and importantly – because we find contrasting patterns for some of our mobility measures, both spatially and temporally.

LS data have been used in the sociology literature to study broad social class mobility; see Buscha and Sturgis (2018) and the earlier work of Platt (2005, 2007), which focuses on ethnicity. This is the first paper, to our knowledge, that exploits the large sample size of the LS, which is an order of magnitude bigger than the traditionally used cohort studies, to generate subnational estimates of intergenerational mobility. A key challenge with the LS is that the census does not record wages or income, and so it is not possible to use this common measure to look at intergenerational mobility. Unlike in many other countries, administrative tax datasets containing earnings or income data in the UK are not sufficiently developed to allow estimates of intergenerational mobility. Instead, this analysis has to use occupation as a measure of intergenerational mobility in the labour market, drawing on a measure of social stratification to rank occupations. The resulting ranking is intuitive and highly correlated with occupation-level wages where comparison is possible. The analysis also studies geographic and temporal shifts in education and homeownership cross-generation mobility.

Our primary contribution with this paper is to generate estimates of intergenerational mobility at a subnational level but, in order to relate to existing work, we are also able to comment on national trends over time. Drawing on a single, much bigger consistent dataset over time and comparing the 1954–1963 birth cohort against the 1974–1983 cohort in 1991 and 2011, occupational mobility is seen to remain broadly the same, consistent with existing evidence based on small surveys. There is some evidence of slightly increased mobility towards the top of the occupation distribution, whilst persistence at the bottom remains starkly constant over time.

In terms of educational mobility, the period under study saw an explosive rise in higher education enrolment and degree acquisition (Blanden and Machin, 2004). This adds extra complications in making clear comparisons over time. Both the total share of parents with degrees and the total share of children with degrees increased in the study period. However, at the start, very few children in the first birth cohort had parents with degrees. This small select group then had a significantly greater chance of obtaining a degree themselves than those whose parents did not hold degrees. Subsequently, the UK's dramatic higher education expansion, starting in the late 1980s and early 1990s, led to many more children of non-degree holders attending university, and so this leads to a high proportionate increase for this group. In terms of the number of times more likely the children of degree holders are to attend university than the children of non-degree holders, educational mobility has risen. However, the percentage-point gap in enrolment probabilities between the two groups actually increased, so by that measure educational mobility fell. Careful definition is needed here to accurately unpack what is going on.

By contrast, the pattern for upward homeownership mobility, broadly defined as the likelihood of owning a home if your parents did not own a home, is clear. This has fallen very dramatically. Homeownership in the UK has become increasingly associated with parental homeownership.

Turning to geographical differences, for most measures there are highly significant, and sizable, differences across areas. Four key insights emerge from this spatial analysis. First, the significant differences in upward occupational mobility across areas prove to be highly persistent over time. Second, and consistent with the existing literature on income mobility, measures of upward mobility – in some cases also reflecting absolute mobility – and relative mobility are positively correlated (especially the case for homeownership mobility). Third, there is a strong link between upward educational mobility and upward occupational mobility, a finding that is consistent with educational attainment being a route to upward occupational mobility, which is an unsurprising result given the known high returns to university education in the UK. Finally, the geography of homeownership mobility is distinctly different from that based on occupation or education. Areas where children of non-homeowners are likely to purchase a home themselves are in fact areas where upward occupational mobility is low. As a very clear example, London is exceptional in its rate of upward occupational mobility, with children from low-occupation fathers being far more likely to reach top occupations than elsewhere, but it scores poorly in terms of homeownership mobility. The broader lesson to take is that homeownership mobility need not align with occupational mobility. Area-level differences in intergenerational mobility are therefore multi-dimensional, and comparisons based on a single metric alone are not necessarily the complete picture.

The structure of the remainder of this paper is as follows. In Section 2, we discuss the data, paying particular attention to the issue of generating numerical scores from categorical occupational classifications. In Section 3, we discuss the related literature and outline methods on how measures of intergenerational mobility from the available data are defined. Section 4 contains the main results, where we present first the national estimates, and then those at the subnational level. We conclude in Section 5 by summarizing the main findings and offering promising directions for future research.

2. Data

The main dataset used in our analysis is the Longitudinal Study of England and Wales (LS). The LS contains linked decennial census and life events data for a 1 percent sample of the population. It contains records on over 500,000 people at any point in time. Records have been linked at each census since the 1971 Census, for people born on one of four selected dates spread throughout the calendar year. These four dates were used to update the sample at the 1981, 1991, 2001, and 2011 censuses. Matching is performed using the National Health Service Central Register (NHSCR). Individual records from each census are matched to the NHSCR, which contains the National Health Service (NHS) number, a unique identifier held by almost all residents of England and Wales.[3] This unique identifier then allows follow-up of individuals over time. Life events data are also linked for LS members, including births to sample mothers, deaths, and cancer registrations. New LS members enter the study through birth and immigration (if they are born on one of the four selected birth dates).

Only LS members can be followed over time, not their extended family. However, in each census, data are provided on all individuals in the LS member's household at the time of census enumeration. Importantly, this means that to obtain matchable data on parental outcomes it is necessary to first observe the LS member at an age where they are living with their parents. The LS member can then be followed over subsequent years to observe their adult outcomes.

The core sample of cohorts covers those individuals aged 8–17 in the 1971, 1981, and 1991 censuses, and who have at least one parent in the same household at this age. This corresponds to birth cohorts of 1954–1963, 1964–1973, and 1974–1983. The adult outcomes for these cohorts are measured two censuses later, when they are aged 28–37.

As a concrete example, the 1974–1983 birth cohort is first observed in 1991 aged 8–17. At this point, data are extracted on parental outcomes; the median age of parents is 41. The 2011 Census can then be drawn upon to measure adult outcomes for the children who are at that point aged 28–37. This age range is selected as it is plausible that by this point, for most individuals, occupational maturity has been reached (Mazumder and Acosta, 2015), and most housing and education decisions have been made. If we were to use children of younger ages at baseline, our estimates would be strongly dictated by life-cycle effects, which are discussed later in the paper. Of course, and especially in the case of housing, it needs to be acknowledged that the ideal age at which to observe outcomes is possibly greater than 28–37, and so robustness checks are reported also using just the older members of the cohort.[4]

[3] NHS numbers are automatically allocated to those born in the UK. However, those who migrate to the UK will only be given an NHS number once they use NHS services. This will be almost all migrants eventually, but some recent arrivals will not yet have needed any health care. The percentage with no NHS number is therefore likely to be small.

[4] Blanden et al. (2021) use cohort study data for the UK to show that estimates of intergenerational homeownership mobility are very similar when cohort members are observed at age 33 and at age 42.

Table 1. Summary statistics by cohort

	1954–1963		1964–1973		1974–1983	
Baseline						
Female (%)	49.9	(50.0)	48.4	(50.0)	48.6	(50.0)
Age at baseline	12.3	(2.9)	12.7	(2.9)	12.5	(2.9)
White (%)	97.3	(16.2)	95.5	(20.7)	91.5	(27.9)
First-generation immigrant (%)	3.4	(18.1)	3.2	(17.6)	3.3	(17.9)
Second-generation immigrant (%)	–	–	9.1	(28.8)	11.3	(31.7)
Missing mother (%)	1.7	(13.0)	3.0	(17.0)	2.1	(14.5)
Missing father (%)	7.0	(25.6)	10.1	(30.1)	14.1	(34.8)
Age of mother at baseline	40.4	(6.8)	39.6	(6.7)	39.4	(5.9)
Age of father at baseline	43.3	(7.3)	42.4	(7.5)	42.2	(6.8)
Parent homeowner (%)	48.8	(50.0)	63.6	(48.1)	75.5	(43.0)
Parents have degree (%)	4.1	(19.7)	8.3	(27.6)	11.3	(31.7)
Mother has degree (%)	0.8	(9.1)	2.9	(16.8)	4.4	(20.5)
Father has degree (%)	4.0	(19.7)	7.9	(27.0)	11.2	(31.5)
Follow-up						
Age at follow-up	32.3	(3.1)	32.7	(3.6)	32.5	(3.1)
Owns house at follow-up (%)	75.6	(43.0)	76.0	(42.7)	66.7	(47.1)
Degree at follow-up (%)	10.6	(30.8)	22.7	(41.9)	39.7	(48.9)
Number of observations	62,609		63,217		50,416	

Notes: Statistics are based on the core sample of individuals aged 8–17 at the baseline year who were successfully followed up two censuses later and matched with at least one parent. Values correspond to means and percentages, depending on the variable. The actual number of observations used to compute average might be lower due to missing data. Standard errors are given in parentheses. Second-generation immigrant not available for first cohort.
Source: ONS LS.

Summary statistics by cohort are presented in Table 1. The baseline and follow-up ages for each cohort are identical. On average, at baseline the individuals in the sample are 12 years old, and at follow up they are 32 years old. One issue to note is that fathers are less likely to be present in the household over time. Given this, where possible, results are also shown using mothers and these prove to be robust. The sample size for the first two cohorts is around 63,000 and falls to 50,000 in the final cohort. This is not a result of differential attrition or other data quality issues, but is driven by changes in the underlying cohort size in the population.

Attrition is low, particularly when compared with cohort studies on which previous UK research on mobility are based. As shown in Table A1 in Online Appendix A, out of the 78,381 individuals born between 1954 and 1963, it is possible to match 73,775, or 94 percent, to at least one parent in the 1971 Census. 65,524, or 84 percent, of these individuals can be traced in the 1991 Census when they are aged 28–37. The Office for National Statistics (ONS) reports that migration, census under-enumeration, and census mismatching can explain attrition (Blackwell et al., 2003).

The final sample is made up of those who can be matched to at least one parent and who can be identified in the 1991 Census. For this cohort, this is 62,609 individuals, or 80 percent of the original sample. For the 1964–1973 birth cohort, the equivalent number is 76 percent and for the 1974–1983 cohort it is 77 percent. Given these low rates of attrition, and because censuses are mandatory, the estimates should be highly representative of the full population of England and Wales for each birth cohort. This is another key advantage of using census data rather than optional surveys. While attrition is low, it is non-random, with older men for example being disproportionately likely to drop out of the sample. Given this, for the national analysis, results are presented using attrition weights.[5]

The empirical analysis of intergenerational mobility focuses on three measures: occupation, homeownership, and education. While earnings or income are not observed, the data contain detailed occupation codes for almost all individuals, including many not in work at the time of the census. Out of those with no occupation given, the majority are long-term unemployed or detached from the labour force. Robustness of the main results to different treatments of the unemployed will be considered later, but the main analysis is based upon all individuals who report an occupation. Across the three cohorts, the number of separate occupations identified ranges from 222 (for the first cohort) to 371 (for the final cohort). Even for the cohort with the fewest separate occupations, only one occupation contains more than 5 percent of the cohort, and the median occupation contains 0.2 percent of the sample. Table A2 in Online Appendix A provides some descriptive statistics on the distribution of occupations across the sample, and shows the top five occupational categories and their share of the observations.

The use of occupations in intergenerational mobility studies has a rich history in sociology. This body of work pre-dates the more recent economics literature. Frequently, occupations are aggregated into broad social class measures, as seen recently applied to LS data in Buscha and Sturgis (2018). However, the aim here is to fully exploit the granularity of the occupational information. While many possible rankings are available, the analysis primarily uses a prominent approach that has been generated with the classifications used in the England and Wales census in mind – the Cambridge Social Interaction and Stratification (CAMSIS) scale (Stewart et al., 1980).[6]

[5]Attrition weights are calculated for each cohort by reweighting the adult sample to match the demographics of the child sample. Match demographics are individual year of birth, gender, and region.
[6]The appendix of Buscha and Sturgis (2018) presents analysis for England and Wales using CAMSIS scores and a different model specification, finding consistent results with what we

The idea behind the CAMSIS project is that occupational rankings can be formed by observing social interactions in the marriage market. Occupations are assigned scalar continuous scores based on the distribution of partner occupations. For example, the distributions of partner occupations for those whose own occupations are doctor or lawyer are close to one another, so these occupations are assigned similar scores. The distribution of partner occupations for farmers is quite different to that of lawyers or doctors, so farmers are assigned a distant score. This is performed separately by gender. A variety of models are used to generate such scores, and full details are given in a recent overview by Lambert and Griffiths (2018).[7] There are alternative approaches to ranking occupations into a continuous measure. These include the International Socio-Economic Index (ISEI) and the Standard International Occupational Prestige Scale (SIOPS). The former scores occupations on their average profiles in terms of income and educational qualifications, whilst the latter uses survey responses from individuals asked to rank the prestige of certain jobs. A review of the approaches is provided by Connelly et al. (2016), whilst Lambert and Bihagen (2014) show that, for UK data, the correlation between CAMSIS and these alternative indices is high (0.84 for ISEI and 0.81for SIOP). It should also be noted that we consider both wage imputation and educational attainment in the paper, which forms the core of the ISEI approach.

The latest set of scores available from the CAMSIS project website are used.[8] These are derived from observed partner occupation distributions in various past censuses and, more recently, the Labour Force Survey. Scores are standardized to have a mean of 50 and a standard deviation of 15 in the population, though ultimately we will apply a rank transformation when using scores. We perform the rank transformation chiefly so that we can draw on the latest rank–rank family of models, which have become the dominant approach to modelling intergenerational mobility in recent years. The resulting occupation ordering is intuitive. For the most recent cohort, the lowest-ranked male occupation is street cleaner, whereas the highest ranked is natural and social science professionals. Each occupation's score is allowed to change over time, reflecting the changing stratification of occupations. While this approach has not been widely used in economics, it represents many advantages in our application and is commonplace in sociology – Savage et al. (2013) provide a well-known recent example. It is possible that the link between occupations and socio-economic status have changed over time,

show here. In particular, they do not use the rank–rank specification but an OLS regression on the raw scores.

[7]Some versions of CAMSIS scores also use information on employment status to form rankings, but here we do not.

[8]See http://www.camsis.stir.ac.uk/index.html.

which would make our results more difficult to interpret. While we cannot entirely rule this out, we stress that this is implicitly assumed in much work using occupations, and it is not possible to assess this empirically using our data. It is also possible that the marriage market has become a more or less informative signal of the socio-economic status of occupations. These concerns are common to all studies using CAMSIS scores.

A key issue when using either occupation or income measures in studies of intergenerational mobility is measurement error. As outlined in Mazumder and Acosta (2015), both income and occupational mobility are overstated when using a single year of occupation compared with a 10-year average centred on mid-career. Unfortunately, our dataset does not allow us to create such an average. Given this, we might expect our occupational mobility estimates to suffer from an upward bias. A second issue is life-cycle effects. As discussed above, we have chosen our sample as best as possible to capture the point at which fathers and children have reached occupational maturity. However, the results of Mazumder and Acosta (2015) suggest that using a later age of measurement for father's earnings might result in lower estimates of occupational mobility, again suggesting a potential upward bias in our estimates. Part of our motivation for using ranks is that these have been shown to be more robust to life-cycle and measurement error issues (Chetty et al., 2014).

As earnings mobility patterns have been a topic of immense debate and policy interest, and are more familiar to economists, in Online Appendix B we discuss a strategy in which external data on wages are integrated into the analysis, through producing a set of occupational wage scores. As the occupations in our wage data do not align well with census occupations, our sample size for the analysis using wages is severely constrained. Therefore, we do not include this analysis in our main results but include a national-level analysis using these data as an appendix for the interested reader. Ultimately, as discussed in Blanden et al. (2013), the CAMSIS score is attempting to measure something fundamentally different from economic resources. Nonetheless, it is highly correlated with wages. For the most recent birth cohort, the correlation between occupational wage percentile rank and CAMSIS score percentile rank is 0.80. Figure A1 in the Online Appendix shows this relationship graphically. While choices of occupation measures make international comparisons difficult, they should be valid for comparisons over time and particularly across geographies, and the discussion here suggests that they are closely related to measures of mobility based on observed incomes.

The census education question changes across years and therefore this restricts the focus to whether the cohort member (and either of their parents) has at least an undergraduate university degree or not, as this is the only way education can be consistently defined. For individuals living with a single parent, the education of only that parent is used. Fortunately, there is extensive evidence to show that this educational distinction is generally the most important one in terms of the labour market benefits to educational attainment. For example, the wage premium for having a degree relative to leaving school with no qualifications in 2017 was 103 percent, whilst the premium for an additional two years of schooling (getting A levels versus no qualifications) was only 28 percent according to calculations from the Labour Force Survey.[9]

Homeownership measures whether the individual owns their home, compared with renting or living in public housing. The analysis does not distinguish between whether the home is owned outright or with a mortgage, nor whether the property is freehold or leasehold. A small number of observations said to be "living rent free" or in "communal establishments" are excluded. These make up fewer than 1 percent of all observations.

While lower-level geographies are available in the census data, to obtain a reasonable sample size within each area, the analysis primarily relies on the European Union's standard geographical classifications (NUTS), which are built on the lower-level local authorities. The main analysis is performed at the NUTS2 level, of which there are 35 in England and Wales. Importantly, because of internal migration, this will not necessarily be the area in which their adult outcomes aged 28–37 are observed. As is common in this literature, "movers" and "non-movers" are not distinguished, although future work intends to explore this dimension.

3. Related research and methods

3.1. Related research

In addition to the UK literature cited in the introduction, this paper adds to the growing set of papers on the geography of opportunity spearheaded by the ground-breaking US work of Chetty et al. (2014). They find, as does this paper, substantial geographic differences in upward mobility. Additionally, they show that high-mobility areas have less residential segregation, less income inequality, better primary schools, greater social capital, and greater family stability. In subsequent work, Chetty and Hendren (2018a, 2018b) generate causal estimates of the impact of neighbourhoods on intergenerational mobility by exploiting family moves from one area to another that occur at different ages for different children within a family. They find that much of the difference in mobility rates across the US are indeed due to causal area effects rather than being driven by sorting. This study is not able to replicate

[9] Own calculations from the 2017 UK Labour Force Survey.

their research design given the challenging data requirements; however, their results are promising for the interpretation of our estimates as being attributed to differences in environment, rather than unobservable differences in individuals.

As this study does, several papers look further back and so are able to investigate the persistence of regional differences in mobility in the US. Focusing on educational attainment rather than income, Card et al. (2018) use the 1940 census data to examine the intergenerational transmission of human capital for children born in the 1920s. They find lower average mobility rates for Black people compared with White people and, consistent with Chetty et al. (2014), wide variation across states and counties. Perhaps the most striking result in their paper is the strong positive correlation between education mobility for the 1920s cohort and income mobility for the 1980–1983 cohort, as measured by Chetty et al. (2014) across the counties of the US. In contrast, Tan (2018) also considered the early 20th century experience using census data, finding that the geography of mobility looks somewhat different to today.

Other studies have investigated geographical differences in intergenerational mobility in Sweden (Heidrich, 2017), Norway (Butikofer et al., 2018), Canada (Corak, 2020), Australia (Deutscher and Mazumder, 2019), and Italy (Acciari et al., 2019). Relative to the tax data used in these papers, the census data studied here offer a longer time series and a wider set of measures, though the sample is smaller, and lacks the detailed income and labour market data many of these studies have access to. Given the current efforts to create longitudinal links in recent US censuses (Alexander et al., 2015), the lessons from our census-based study ought to be particularly relevant for the US setting, where it is likely that similar data will be made available soon.

3.2. Methods

Here, we briefly discuss methods used to generate measures of intergenerational mobility across the three components. These are standard in the literature and so the interested reader is referred to more in-depth discussions where appropriate.

The main measure for intergenerational occupational mobility is the rank–rank relationship first used by Dahl and DeLeire (2008) and popularized by Chetty et al. (2014). Ranks are calculated based on the CAMSIS scores discussed in Section 2, rather than using more common measures of income or earnings. Each individual is assigned a percentile rank based on their position in the national occupational distribution for their cohort, and ranks for each parent are calculated equivalently. As CAMSIS scores are not comparable between men and women, this is calculated separately for each gender. For example, a rank of 70 assigned to a father in a birth cohort means that 70 percent of fathers work in occupations that are ranked lower than this particular father's occupation. In the remainder of the text, these are referred to as occupational ranks. On the parent side, a father's occupation is used as it has fewer missing values, particularly for early cohorts. For the national estimates, as shown below, the results are similar if a mother's occupation is used instead.

For each cohort, a linear model relating child rank (R_i) to parent rank (P_i) is

$$R_i = \alpha + \beta P_i + u_i. \tag{1}$$

In equation (1), the intercept α reflects the average rank of those with parents at the very bottom of the occupation rank distribution, whereas the slope β reflects the average increase in child rank associated with a one-rank increase in parent rank. Higher values of β reflect lower rates of relative mobility. In the results, the rank–rank intercept estimates are denoted by $\hat{\alpha}^{OCC}$ and the slope estimates by $\hat{\beta}^{OCC}$. A key advantage of the rank–rank model is that it suits subnational analyses well. The regression is run separately on individuals in each geographical area, and, consistent with the literature, ranks are always being defined relative to the national distribution.

A commonly used measure of absolute upward mobility at the subnational level is the expected rank, according to the above regression, of children whose parents are at the 25th percentile of the distribution. Across areas in the spatial analysis, one such metric is the average occupational rank of those whose fathers are close to the bottom of the occupational distribution. For comparison, the expected rank of those whose parents are at the 75th percentile can be computed. These expected ranks are denoted by $\hat{R}^{25}$ and $\hat{R}^{75}$, respectively.[10]

To complement the rank–rank analysis, transition matrices are also presented at the national level. These impose less structure on the modelled relationship than the rank–rank model. There are many different ways of quantifying dependence in transition matrices (Jäntti and Jenkins, 2013). The main measure used here is the immobility ratio (IR), which is given by the proportion of all observations found on the leading diagonal of the transition matrix. A particular focus is placed on the outcomes of cohort members who start in the bottom group, as is typical in the literature.

Degree attainment and homeownership are binary variables for both parent and child. At the national level, transition matrices are presented, though linear probability and logit models are also estimated. Intercept estimates from these models, now given by $\hat{\alpha}^{DEG}_{LPM}$, $\hat{\alpha}^{DEG}_{LOG}$, $\hat{\alpha}^{HOME}_{LPM}$, and $\hat{\alpha}^{HOME}_{LOG}$ can be thought of as indicators of upward mobility, as they reflect the probability of individuals

[10]Consistent with the literature, control variables are not included.

who start in the "low" state moving to the "high" state. The slope estimates, again denoted by $\hat{\beta}$ with the analogous subscripts and superscripts, reflect relative mobility. When the base levels (e.g., the probability of obtaining a degree for those whose parents did not have a degree) are close to zero, the two can give quite different results. The logit coefficients reported are interpreted as the effect of a child having a parent in the high state (e.g., with a degree) on the log-odds of the child reaching the same state. If a group initially has a probability of close to zero of obtaining some outcome, even a small percentage change in that probability will correspond to a large coefficient. In practice for our setting, this matters for comparisons of educational mobility, but not for homeownership. For subnational comparisons, the linear probability model estimates are used, which in that setting are almost identical to the logit estimates.

4. Results

4.1. National results

4.1.1. Occupational ranking. The starting point is to show aggregate results for the various measures of mobility and, where possible, to compare them with extant estimates in the literature, both for the UK and for other countries. After this, results are presented about the key contribution of this paper – the changing geography of mobility in England and Wales. Whilst it is possible to present results for each birth year, we do not do so here as the sample sizes become quite small, leading to imprecise, noisy estimates.

Figure 1 shows binned scatter plots of child occupation rank on father occupation rank for each cohort. Each point represents the average occupational rank of fathers and children within one of 40 quantiles of the distribution of the father occupation rank. There is a strong positive relationship that is approximately linear for all three cohorts, perhaps with a slight convexity. Interestingly, unlike in much work on income mobility, we do not find an S-shaped relationship, with the gradient steepening at the top and bottom of the distribution of the father occupation rank. The same is true if we use occupation–wages rather than CAMSIS scores.

The graphical evidence presented in Figure 1 presents a justification for the linear specification adopted in the rank–rank regressions, the results of which are given in Table 2. Looking at the first column, the slope coefficient estimate for the 1954–1963 birth cohort is 0.36, meaning that a one-unit increase in father occupation rank corresponds to an increase of 0.36 in own occupation rank. For all periods, the relationship is stronger for men. The estimated slope in the final birth cohort is slightly lower at 0.32, suggesting a small increase in mobility. This is consistent with the previous work of Buscha and Sturgis (2018) as well as the longer-term historical evidence

Figure 1. Occupation rank–rank relationship

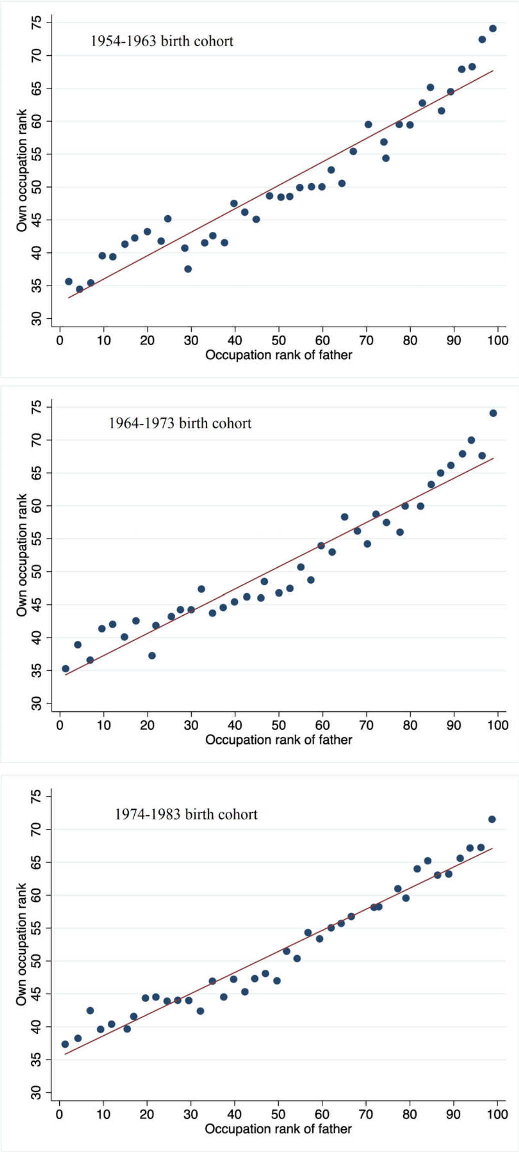

Notes: Binned scatter plots of own occupation rank on occupation rank of father for each birth cohort. Points displayed are values of own and father occupation rank in 40 quantiles of occupation rank of father. Cohorts aged 8–17 at baseline (parent outcome measured), and 28–37 at follow-up (child outcome measured). Ranks are determined by CAMSIS score and calculated separately within gender. No controls are included.
Source: ONS LS.

Table 2. Occupation rank–rank estimates

	(1)	(2)	(3)
1954–1963			
$\hat{\beta}^{OCC}$	0.357	0.373	0.338
	(0.004)	(0.006)	(0.006)
$\hat{\alpha}^{OCC}$	32.45	31.65	33.35
	(0.238)	(0.324)	(0.351)
Gender	Both	Male	Female
N	52,360	27,655	24,705
1964–1973			
$\hat{\beta}^{OCC}$	0.337	0.359	0.315
	(0.004)	(0.006)	(0.006)
$\hat{\alpha}^{OCC}$	33.92	32.59	35.19
	(0.240)	(0.341)	(0.338)
Gender	Both	Male	Female
N	52,317	25,645	26,672
1974–1983			
$\hat{\beta}^{OCC}$	0.321	0.329	0.312
	(0.005)	(0.007)	(0.007)
$\hat{\alpha}^{OCC}$	35.43	34.74	36.10
	(0.272)	(0.387)	(0.382)
Gender	Both	Male	Female
N	40,694	20,135	20,559

Notes: Estimates from regression of child occupation rank on father's occupation rank. Ranks determined by CAMSIS scores and calculated within gender. No controls are included. Cohorts aged 8–17 at baseline (parent outcome measured), and 28–37 at follow-up (child outcome measured). Standard errors are given in parentheses. Change from earliest to latest cohort significant at the 1 percent level for all groups.
Source: ONS LS.

shown in Lambert et al. (2007). The result holds for both genders combined and each gender individually. Table A3 in Online Appendix A shows that this is robust to attrition weights, dropping the unemployed and using mother rather than father occupational rank.[11] The estimate of 0.32 sits just below the US income rank–rank estimate of 0.34 reported by Chetty et al. (2014), and above the most comparable UK earnings estimate of 0.27 found in Rohenkohl (2019).

While the parametric rank–rank model provides a single tractable estimate of mobility, a more detailed picture is shown in Table 3, which presents

[11]Confirming that the results are similar when we drop the unemployed is important because, as a referee pointed out to us, within a given occupational category, children of high-income parents are less likely to experience unemployment than children of low-income parents. This would cause a downward bias in the rank–rank slope and an upward bias in the intercept of the occupation measure used in this paper.

Table 3. Occupation transition matrices

Parent	Child				
	1	2	3	4	5
1954–1963					
1	32.4%	25.1%	18.2%	15.6%	8.6%
2	27.7%	25.2%	19.7%	17.7%	9.6%
3	19.7%	22.2%	23.0%	21.1%	13.8%
4	13.3%	16.8%	20.6%	27.7%	21.6%
5	6.8%	11.1%	16.7%	28.8%	36.4%
Sample size			52,360		
1964–1973					
1	33.5%	24.2%	18.0%	15.1%	9.1%
2	24.7%	25.7%	19.7%	17.5%	12.3%
3	20.6%	24.3%	20.3%	19.6%	15.0%
4	14.1%	18.2%	20.6%	24.8%	22.1%
5	7.8%	12.3%	16.7%	28.2%	34.7%
Sample size			52,317		
1974–1983					
1	32.7%	24.2%	19.2%	14.0%	9.7%
2	25.5%	25.2%	20.2%	16.5%	12.3%
3	20.2%	22.3%	22.2%	18.7%	16.2%
4	13.1%	18.8%	20.7%	23.6%	23.7%
5	8.5%	12.6%	18.2%	26.3%	34.2%
Sample size			40,694		

Notes: Quintiles of occupation ranks within each cohort. Ranks determined by CAMSIS scores and calculated within gender. Father's score used for parent quintile. Each row sums to 100 percent within each cohort.
Source: ONS LS.

quintile transition matrices. The IR is 0.29 in the first cohort, meaning that 29 percent of the mass falls on the diagonal in the first cohort. If there were no dependence, this would be 20 percent. The IR falls slightly to 0.28 in the final cohort. In the final cohort, individuals are marginally less likely to end up in the same occupation rank quintile as their father, consistent with the rank–rank analysis. However, inspecting the outcomes of those at the bottom of the distribution, the probability of those who start in the bottom quintile remaining in the bottom quintile is broadly unchanged, moving from 32.4 percent to 32.7 percent. Persistence at the bottom is remarkably stable over time. The same is true for those in the second-lowest quintile.

In contrast to the bottom, there is an increase in churn at the top. Top-quintile persistence has fallen from 36.4 percent to 34.2 percent, with those from all other four quintile groups now more likely to reach the top. The probability of a child from the bottom quintile of occupations reaching the top quintile is 8.6 percent for the first cohort, increasing to 9.7 percent in the final cohort. In terms of cross-country comparisons, using tax data, Chetty et al. (2014) report a value of 7.5 percent for the bottom-to-top quintile transition probability in the US, whilst figures for Denmark (Boserup et al., 2013) and Canada (Corak and Heisz, 1999) are 11.7 percent and 13.4 percent, respectively. Whilst it is necessary to exercise caution in comparing transition matrices based on household income or labour market earnings with the estimates based on occupational rankings, the bottom-to-top transition probability sits between the low rates of the US and the higher rates in Denmark and Canada. This is consistent with previous cross-country comparisons of mobility rates (such as Blanden et al., 2005).

To summarize the national occupational mobility results, using occupational ranks derived from CAMSIS scores, mobility has been broadly flat across birth cohorts. There is some evidence of more churn at the top of the distribution, but persistence at the bottom is unchanged. As outlined in Online Appendix B, using occupation wages rather than CAMSIS scores, there is little overall change. Interestingly, and consistent with other work, there is a one-time fall in mobility between the 1954–1963 and 1964–1973 birth cohorts. The fact that this is not detected in our CAMSIS score analysis is interesting and consistent with previous comparisons of occupation and wage-based measures, as discussed in Blanden et al. (2013).

4.1.2. Educational attainment. Table 4 reports mobility estimates for degree attainment using simple transition matrices. We also report both logit and linear probability model (LPM) specifications for the binary outcome in Table A5. This period saw a substantial expansion of higher education in the UK, seen in rising probabilities of obtaining a degree for both parents and children. For the first cohort, only 9 percent of children obtained a degree if neither of their parents had graduated, compared to 49 percent for those with a graduate parent. By the last cohort, these figures had risen to 35 percent and 80 percent, respectively. For this outcome, there is a difference between the logit and LPM estimates. Between the first and second cohorts, the logit slope falls from 2.295 to 2.048, and this change is highly significant, corresponding to an increase in mobility. The LPM slope, however, increases from 0.40 to 0.45. Between the second and third cohorts, neither model shows much change.

As already remarked, the divergence between the two measures in the first period arises because the probability of obtaining a degree for children of non-degree parents in the earliest period was exceptionally low. Therefore, even a small increase in this probability generates a large change in log-odds for this group relative to those whose parents had degrees, thereby bringing about a fall in the logit coefficient. On aggregate, for the first cohort, only 4 percent of parents had a degree and 11 percent of the children. For the second, these figures rose to 8 percent and 23 percent, and for the final cohort, 11 percent and 40 percent, respectively.

Table 4. Education transition matrices

Parent	Child		
	Degree	No degree	Parent sample size
1954–1963			
Degree	49.1%	50.9%	2,542
No degree	9.0%	91.0%	60,067
Child sample size	6,634	55,975	
1964–1973			
Degree	64.4%	35.6%	5,244
No degree	18.9%	81.1%	57,845
Child sample size	14,332	48,757	
1974–1983			
Degree	79.5%	20.5%	5,701
No degree	34.7%	65.3%	44,679
Child sample size	20,018	30,362	

Notes: Education is a binary measure of attainment of an undergraduate degree qualification. Each row sums to 100 percent within each cohort.
Source: ONS LS.

Factoring this in, the early part of the large expansion in UK higher education, according to the rise in the LPM measure coupled with the fact that so few parents had degrees in the first period, therefore effectively acted to increase the link between parent and child university education, in that it almost guaranteed a university education for those whose parents had also had one. Furthermore, both measures are in agreement that between the 1964–1973 and 1974–1983 birth cohorts, there was no change to degree mobility despite a near doubling of the probability of university attendance. Children of degree holders in the most recent cohort now have an 80 percent probability of going to university, compared to a probability of 35 percent for those whose parents did not attend university. The gap between the raw probability of a child getting a degree if at least one of the parents did have a degree compared with if they did not actually rose from 0.40 to 0.45, but the ratio of the probabilities between the two groups declined. In Table A6 in the Online Appendix, this is shown to be robust to attrition weights and restricting to the older sample members.

4.1.3. Homeownership. Table 5 shows the transition matrix results for homeownership, with corresponding logit and LPM results reported in Table A7. There has been a large and statistically significant fall in intergenerational mobility across the cohorts – homeownership among children has increasingly depended on the homeownership of their parents. In the first cohort, children were 19 percentage points more likely to own their home if their parents did,

Table 5. Homeownership transition matrices

Parent	Child		
	Owner	Renter	Parent sample size
1954–1963			
Owner	85.1%	14.9%	31,337
Renter	66.5%	33.5%	34,142
Child sample size	48,763	16,716	
1964–1973			
Owner	83.9%	16.1%	41,030
Renter	62.2%	37.8%	24,642
Child sample size	48,553	17,119	
1974–1983			
Owner	74.3%	25.7%	38,343
Renter	43.3%	56.7%	13,391
Child sample size	34,092	17,642	

Notes: Homeownership is a binary measure of whether at least one parent owned their own home. Each row sums to 100 percent within each cohort.
Source: ONS LS.

compared with those whose parents rented. For the final cohort, this gap had risen to 31 percentage points.

Inspecting the summary statistics in Table 1, those born in the late 1950s were significantly more likely to own their homes than their parents, but for those born in the late 1970s the opposite is true, with homeownership rates falling below 70 percent. The mobility estimates reveal that this drop was disproportionately felt by those whose parents did not own a home. For the first cohort, 67 percent owned their own home even if their parents did not, whilst this had dropped hugely to 43 percent for the last cohort. For those whose parents owned a home, the drop was from 85 percent to 74 percent. While homeownership rates have fallen for both groups, they have fallen most rapidly for those whose parents did not own their home. Again, in Table A8, these conclusions are shown to be robust to attrition weights and restricting to the older sample members. Showing that the results hold for older sample members gives us confidence that our results are not being driven by changing ages of first homeownership.

Blanden and Machin (2017) also provide evidence on homeownership mobility using the two cohort studies referred to above. Using the same empirical approach to that used here, they find that the LPM slope coefficient goes from 0.140 for the 1958 cohort to 0.217 for the 1970 cohort. This is a somewhat larger change in the estimates for the first two cohorts (though their estimates are for when the children are aged 42), but again points to a clear decline in mobility along this dimension. The evidence here covers an additional later cohort, showing that this decline accelerates.

4.2. Subnational results: the geography of mobility

Moving to geographic intergenerational mobility, subnational NUTS2-level estimates are next considered. These areas all have a population of at least 800,000 individuals. In the data, the average sample size is 1,689 observations per area, ranging from a minimum of 306 to a maximum of 3,903. The full set of NUTS2-level estimates are given in Table A12 in Online Appendix A, and are plotted along with confidence intervals in Figures A2, A3, and A4.

There are some important considerations when interpreting these spatial results. The regressions do not control for other characteristics which may differ across individuals in areas. When interpreting occupational mobility estimates, for example, it is important to bear in mind that there are likely to be other features of parents that could explain differences (e.g., immigrant status). The aim here is to generate summaries of the potentially complex mechanisms underlying the associations between parents and children, and to discuss how these summaries relate to one another and how they change over time. This can be seen as a necessary first step that can prompt future work on understanding why there are differences across areas, whether this is due to unobserved features of parents, or due to policy, peers, or other differences in environment. For income mobility in the US, Chetty and Hendren (2018a, 2018b) followed up their earlier, highly influential descriptive work by convincingly demonstrating that much of the regional differences are the result of causal area effects rather than selection.

4.2.1. Occupational ranking. First, consider occupational mobility. Figure 2 shows a choropleth map of the NUTS2 estimates for upward occupational mobility, as given by the expected occupational rank of children of fathers at the 25th percentile of the national occupation distribution ($\hat{R}^{25}$). We show maps for the first cohort, born 1954–1963 (Figure 2(a)) and for the most recent cohort, born 1974–1983 (Figure 2(b)) so that comparisons can be made over the full sample period. The equivalent figure for the middle cohort, born 1964–1973 is shown in Figure A5 in the Online Appendix. Darker shades correspond to greater upward mobility in the maps, and a strong geographical pattern is evident in both graphs.

More specifically, in the first map, a striking contrast can be seen between the five NUTS2 areas of London, towards the south-east corner of the map, and everywhere else in the country. The area with the highest upward occupational mobility for the 1954–1963 birth cohort, Inner West London, has an estimate of 49.6. This means that children from the 25th percentile occupation rank on average end up at the median occupation rank for the country. Children born to low-ranked occupation fathers in Inner West London suffered no occupational disadvantage relative to the national average. In terms of transition probabilities, children in Inner West London starting in the bottom

Figure 2(a). Upward occupational mobility, 1954–1963 birth cohort

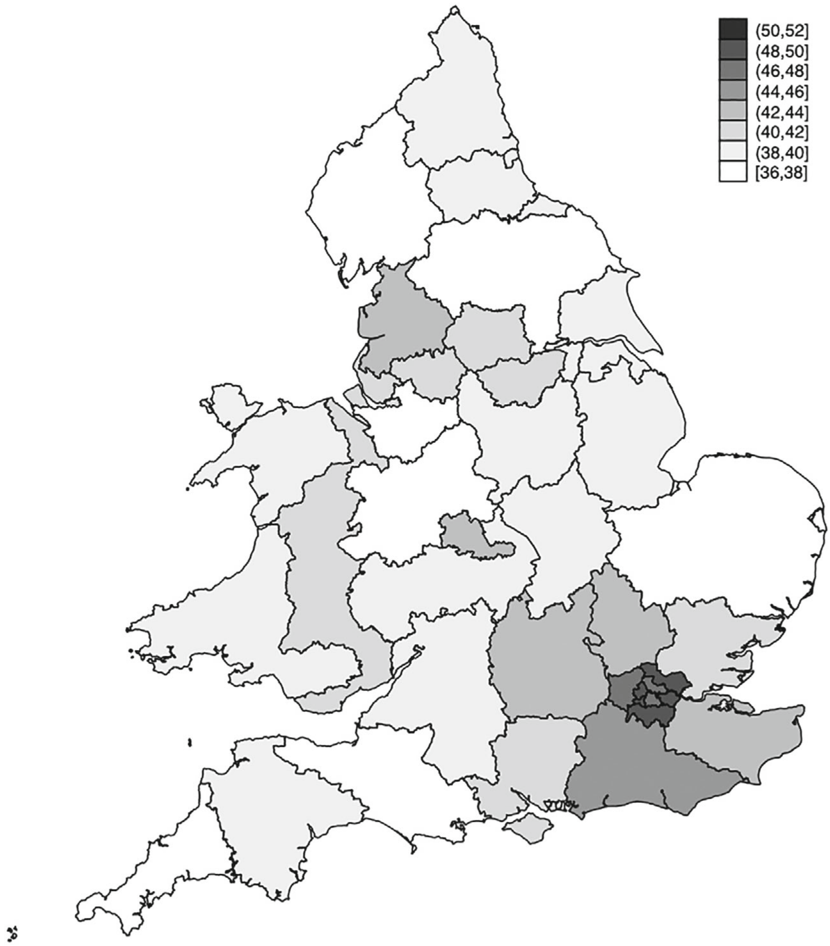

Notes: Maps of R^{25} for 35 NUTS2 areas. Deeper shades correspond to higher values.
Source: ONS LS.

tercile of the occupation distribution had a 28 percent chance of reaching the top tercile. The lowest scoring area, East Anglia, has an $\hat{R}^{25}$ of 36.9, and children growing up in the bottom tercile there had only a 16 percent chance of reaching the top tercile.

Figure 2(b) shows a similar geographical pattern to Figure 2(a), with London still standing out as the most upwardly mobile area in terms of occupation. The correlation between the 1954–1963 and the 1974–1983 estimates of $\hat{R}^{25}$ is 0.69, as shown in Table A9. This corresponds to a high

Figure 2(b). Upward occupational mobility, 1974–1983 birth cohort

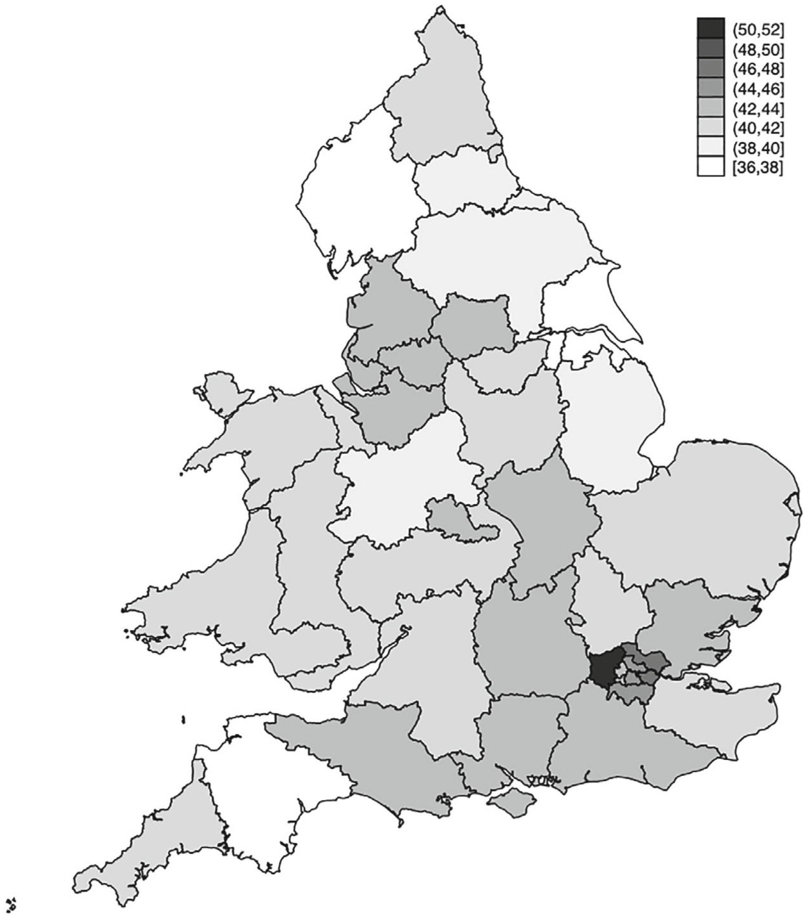

degree of persistence in mobility, consistent with the US evidence discussed previously. The range of estimates is slightly greater than the early cohort, with an $\hat{R}^{25}$ for Cumbria of 36.0 and an $\hat{R}^{25}$ for London (Outer W/NW) of 51.0, though this is driven by the latter area being an outlier. Children growing up in Outer West and North West London in the bottom tercile of occupations had a 32 percent chance of reaching the top tercile, relative to a 12 percent chance in Cumbria. Children of low-occupation fathers in the suburbs of London are more than twice as likely to reach the top occupations than children of similar fathers in the North West.

The choropleth maps are suggestive of some convergence between the rest of the country and London. Colour bins in these maps are fixed intervals, so the swathe of middle-shaded areas in Figure 2(b) suggests that there has been some catch up of areas outside London, particularly in the South. Outer West and North West London, however, remain significant outliers and there is a clear overall North–South divide.

These results echo those from Chetty et al. (2014), finding that mobility tends to be higher in suburban areas outside of major population centres. Interestingly, Inner London was upwardly mobile in the mid-20th century, but mobility had declined substantially by the 1970s and 1980s birth cohorts. The $\hat{R}^{25}$ for the Inner West London 1954–1963 birth cohort was 49.6, but for the 1974–1983 birth cohort it was 43.4. This is consistent with the results of Friedman and Macmillan (2017), who also find low mobility in Inner London for recent cohorts, though their analysis is based on current residence rather than place of birth.

In terms of relative mobility $\hat{\beta}^{OCC}$ (not mapped), there is quite significant spatial variation, with estimates ranging from 0.26 in Essex to 0.42 in Inner West London for the most recent cohort. There is almost no within-area persistence in this rank-based measure of relative mobility. These are estimated less precisely than the intercepts and it is not possible in statistical terms to reject the null hypothesis of equality. Nonetheless, consistent with Chetty et al. (2014), we find a strong negative correlation between the intercept and slope estimates (lower meaning more relative mobility). To the extent that a positive intercept reflects upward mobility, and absolute upward mobility in the case of the education and homeownership variables considered below, which are not based on within-cohort ranks, this is in line with measures of absolute and relative mobility tending to be spatially positively correlated. Table A10 shows that for the most recent cohort, the correlation between $\hat{\beta}^{OCC}$ and intercept $\hat{\alpha}^{OCC}$ is −0.57, and between $\hat{\beta}^{OCC}$ and $\hat{R}^{25}$ is −0.34.

Table A11 shows that in terms of changes in mobility estimates over time this relation is particularly strong, increasing to −0.90. Areas where children of low-occupation fathers do well are those for which relative occupational mobility is highest. Returning to the correlations for the most recent cohort, while the correlation of $\hat{R}^{25}$ and $\hat{R}^{75}$ is very high at 0.8, the variance and range of $\hat{R}^{75}$ are both marginally lower than those for $\hat{R}^{25}$. This can be read as evidence that areas matter more for those towards the bottom of the occupation distribution. This will be returned to when discussing homeownership mobility, where the evidence is far stronger both empirically and conceptually, because the measure can be thought of as showing absolute upward mobility for a positive intercept coefficient.[12]

[12]Whether the presence of a negative spatial correlation between the estimated α intercepts and β slopes reflects a positive relation between absolute and relative mobility varies for the

4.2.2. Educational attainment. Moving to estimates of education mobility, geographical patterns of upward degree mobility for the most recent cohort are shown in Figure 3.[13] It shows the LPM intercept $\hat{\alpha}^{DEG}_{LPM}$, which has the convenient interpretation of being the share of those whose parents do not have degrees obtaining degrees themselves. Logit models give almost identical estimates, so these are not presented. Whilst the specifications do not include controls, estimates are essentially unchanged by their inclusion. Again, Outer West and North West London stand out as being exceptionally upwardly mobile. For those born between 1974 and 1983 in these areas, 45 percent of children of non-degree holders went on to obtain degrees themselves. The figure for the lowest-ranked area on this measure, South Yorkshire, is 29 percent. The null hypothesis of equality of $\hat{\alpha}^{DEG}_{LPM}$ estimates can be rejected at the 1 percent level.

A comparison of the most recent cohort with the 1954–1963 birth cohort shows that there is again strong persistence, with a correlation of 0.71, similar to the occupation results. This is perhaps surprising, given the large expansion in higher education that came over this period. Areas that were successful in seeing children of non-university-educated parents graduate from university in the early cohort were also those that were successful at doing so in the final cohort. In terms of slope coefficient $\hat{\beta}^{DEG}_{LPM}$, in contrast to the occupation results, there are now statistically significant differences, shown in Figure 4. Inner London stands out here as having a particularly high $\hat{\beta}^{DEG}_{LPM}$ estimate. In Inner West London, children of parents with degrees have a 55 percentage point higher chance of going to university than children of parents without degrees.

Comparing occupation and education mobility, there is a strong relationship between $\hat{\alpha}^{DEG}_{LPM}$ and $\hat{\alpha}^{OCC}$. Table A10 shows a correlation of 0.46. This is unsurprising, given that degrees yield high wage returns in the UK, and are necessary to reach many of the top occupations. Table A11 shows that the relationship is also true over time. Comparing changes in upward occupation mobility and education mobility over time between the first and last cohorts, areas that became more upwardly mobile in terms of education also saw greater increases in upward occupation mobility, with a correlation of 0.37.

---

measures considered here and on whether rank–rank models are estimated. For the occupation rank-based measures, the relation is at least in part a mechanical one due to the use of ranks. But for the education and homeownership results to come, a positive intercept can be seen as absolute upward mobility. To further shed light on this, later we also show how the occupation equation intercepts correlate with the intercepts from the non-rank-based specification estimates for these two outcomes.

[13] We show the results for the 1974–1983 cohort to illustrate the most recent evidence on geographical patterns of intergenerational mobility. Equivalent figures for the first two cohorts for both education and homeownership are given in Figures A6 and A7 in the Online Appendix.

Figure 3. Upward degree mobility, 1974–1983 birth cohort

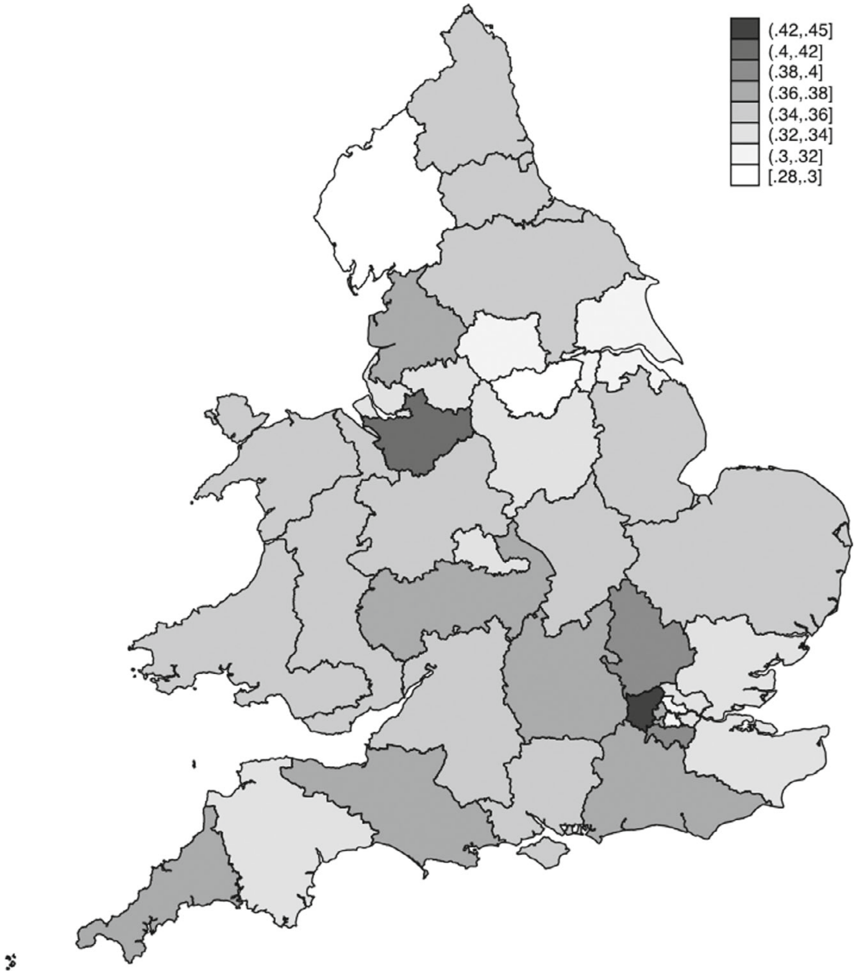

Notes: Map of $\hat{\alpha}^{DEG}_{LPM}$ for 35 NUTS2 areas. Deeper shades correspond to higher values.
Source: ONS LS.

As discussed in Chetty et al. (2014), this relationship suggests that upward mobility in the labour market might be more about childhood environment rather than local labour market conditions. The correlation between relative education and occupation mobility is positive, but somewhat weaker. In addition, the strong relationship between upward education and occupation mobility suggests that comparisons based on upward education mobility might be reasonably accurate proxies for upward occupation or wage mobility.

Figure 4. Relative degree mobility, 1974–1983 birth cohort

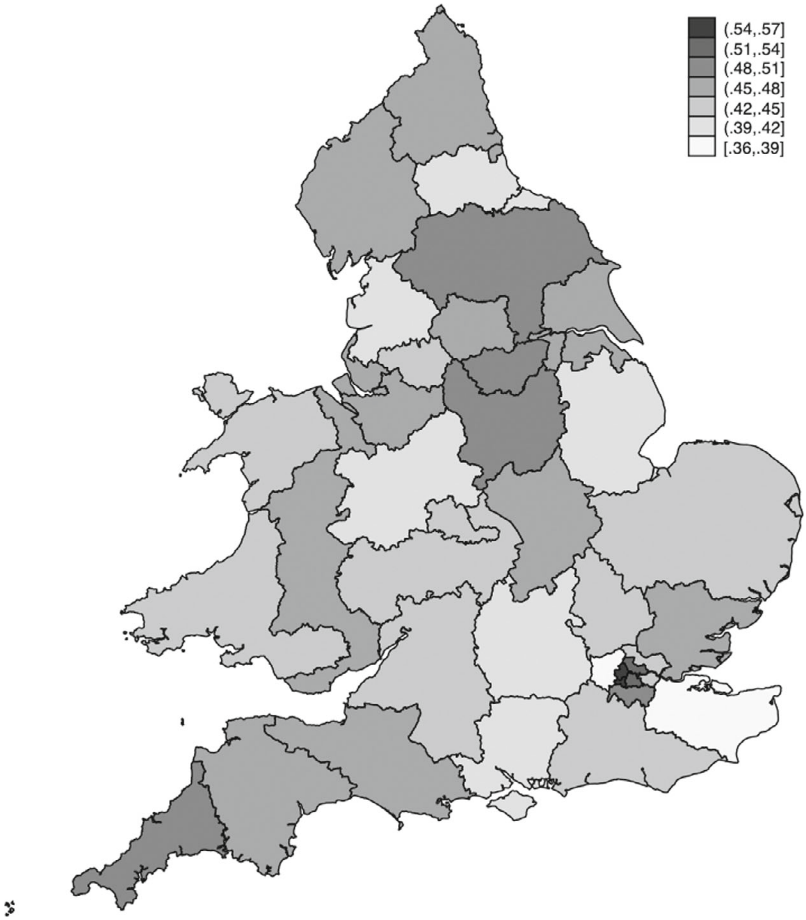

Notes: Map of $\hat{\beta}^{DEG}_{LPM}$ for 35 NUTS2 areas. Deeper shades correspond to higher values.
Source: ONS LS.

Given the less stringent data requirements for generating education mobility measures, this is potentially a very useful result.[14]

4.2.3. Homeownership. Having discussed occupation and education mobility and demonstrated a strong link between the two, the geography

[14] Blanden (2013) shows that this correlation between occupational and educational mobility is also a feature of cross-country data.

Figure 5. Upward homeownership mobility, 1974–1983 birth cohort

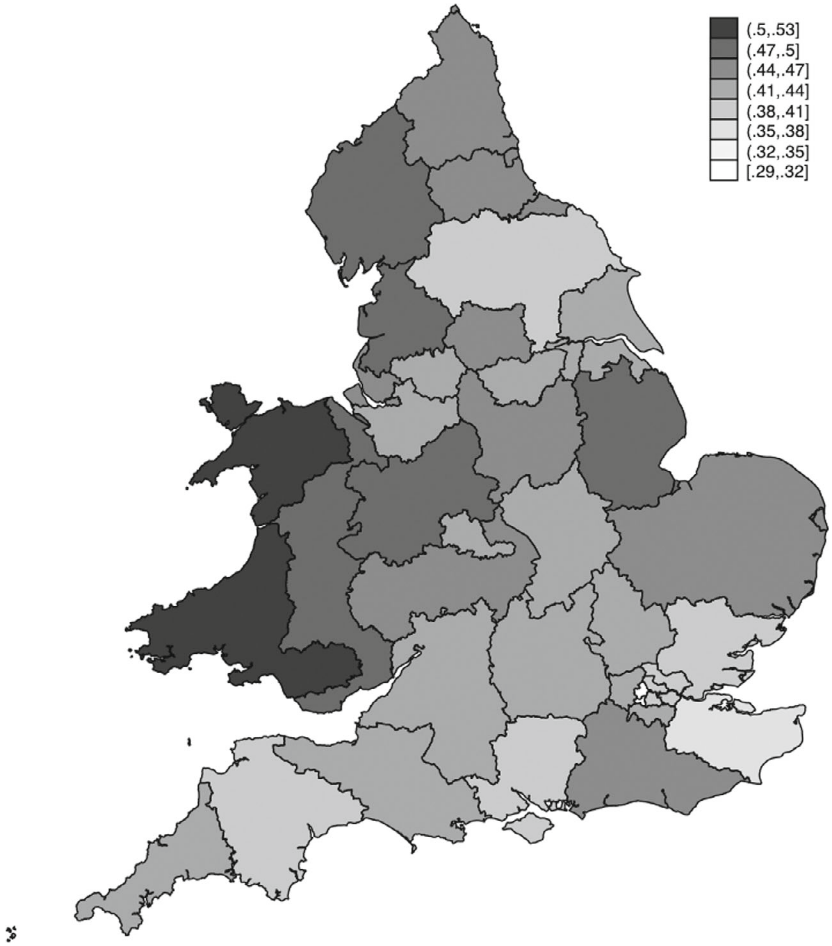

Notes: Map of $\hat{\alpha}_{LPM}^{HOME}$ for 35 NUTS2 areas. Deeper shades correspond to higher values.
Source: ONS LS.

of homeownership mobility can be considered. Figure 5 shows $\hat{\alpha}_{LPM}^{HOME}$ for the 1974–1983 birth cohort. There is very substantial variation, with Inner West London seeing a low estimate of 29 percent, and West Wales at the other end of the distribution with 52 percent. Areas in the North and particularly in Wales have high rates of upward homeownership mobility, whereas the London areas all receive low estimates.

Figure 6 shows the map of the homeownership slope estimates, $\hat{\beta}_{LPM}^{HOME}$. The map is almost the inverse of $\hat{\alpha}_{LPM}^{HOME}$, reflecting an exceptionally strong

Figure 6. Relative homeownership mobility, 1974–1983 birth cohort

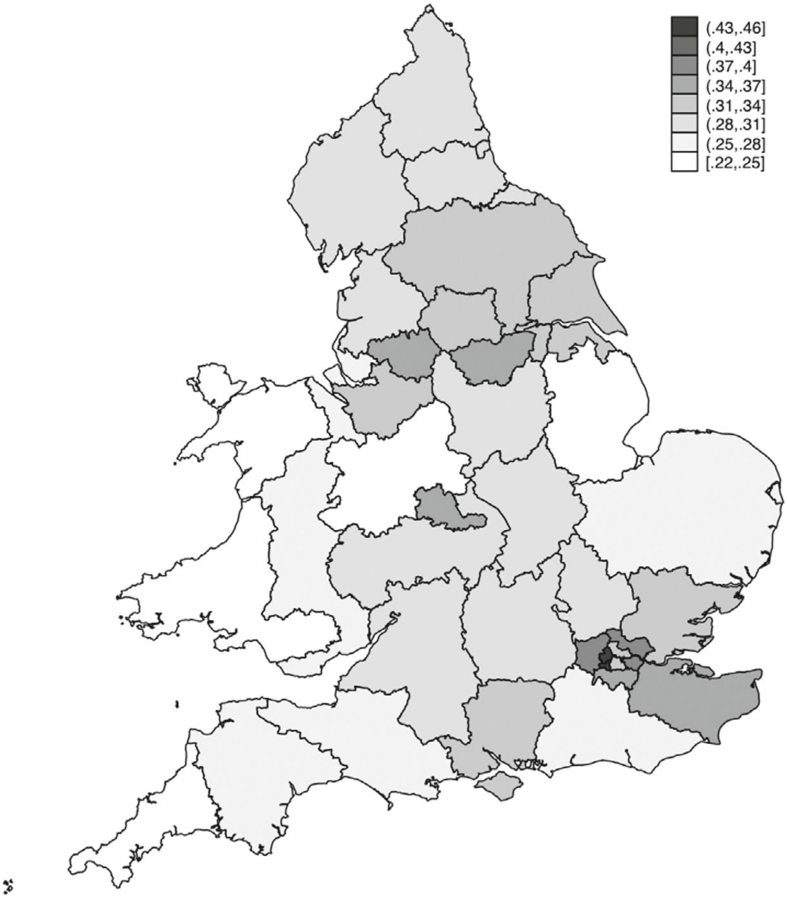

Notes: Map of $\hat{\beta}_{LPM}^{HOME}$ for 35 NUTS2 areas. Deeper shades correspond to higher values.
Source: ONS LS.

positive relationship between upward absolute homeownership mobility and relative homeownership mobility. In Inner West London, children whose parents own their home are 44 percentage points more likely to own a home themselves than those whose parents do not. In West Wales, the gap is only 22 percentage points. The probability of owning a home conditional on parental homeownership $\left(\hat{\alpha}_{LPM}^{HOME} + \hat{\beta}_{LPM}^{HOME}\right)$ has a far lower range than $\hat{\alpha}_{LPM}^{HOME}$, from 66 percent in Cornwall to 80 percent in Outer West/North West London. In other words, geography matters more for those whose parents do not own a home than for those whose parents do own a home.

The fact that geography matters more for non-homeowners than homeowners is, as far as we know, a novel result. While it is not possible with the available data to directly inspect the mechanisms, many of the areas with low upward mobility in homeownership have high house prices (notably around London), having experienced rapid house price growth in the 1990s and 2000s. In these areas, assuming no migration, children of homeowners might be able to rely on family wealth to purchase a home, whereas children of non-homeowners are unable to get a foot on the housing ladder. The structure of the housing market in London differs to other regions, with more renters in London than other NUTS2 areas. A limitation of the analysis here and elsewhere in the paper is that we do not have the sample size to investigate other cities in England and Wales, all of which are significantly smaller than London.

Area-level estimates of homeownership mobility are less persistent over time than shown above for occupation and education, as shown in Table A9. This is again consistent with a strong role of house prices. The house price growth experienced in the 1990s and 2000s was disproportionately felt in certain areas and led to a significant change in the geography of house prices, with many ex-industrial areas seeing house price falls over the same period.

Table A10 shows that the correlation between $\hat{\alpha}^{HOME}_{LPM}$ and $\hat{\alpha}^{OCC}$ is small and negative at −0.157. As homeownership is a key component of wealth, this is suggestive that regional differences in either occupation or education mobility might not reflect differences in wealth mobility. Many of the correlations between housing mobility and other measures are driven by London, which tends to be located at either end of the range of estimates. London might be a fantastic place for less-privileged children to grow up in terms of reaching high-status occupations themselves, but they have little chance of getting a foot on the housing ladder. This means that we ought to be cautious before claiming that certain areas are mobility hotspots in employment-based measures, such as occupation or earnings. Housing is an important part of wealth accumulation and has broader non-financial benefits, and these mobility hotspots might be precisely those areas where access to the housing market is limited for those of poorer backgrounds.

5. Discussion and conclusion

The focus of effort in this paper has been to use the best available data to rigorously generate measures of intergenerational mobility for multiple cohorts at the national level and across geographies of England and Wales over time. We present new estimates of trends and geographical differences in rates of intergenerational mobility in England and Wales over the mid to late 20th century. Comparing a cohort born in 1954–1963 to the cohort born 20 years later, national-level results point to little overall change in occupation mobility, with some increases in churn at the top of the distribution, but little change at the bottom. In terms of educational mobility, the huge expansion in higher education did lead to more first-in-family degree holders but, particularly between the 1964–1973 and 1974–1983 birth cohorts, it did not clearly improve educational mobility. While many children of non-degree holders were able to study at university, children of degree holders were still far more likely to attend university in the final cohort, and by some measures this gap has increased. In terms of homeownership mobility, the picture is very stark. The probability of owning a home if your parents did not own a home has plummeted over the period. Homeownership now is far more dependent on whether your parents owned a home or not.

The geographical analysis, the most detailed yet for England and Wales, points to significant differences in intergenerational mobility across areas. Here, there are four main academic contributions. The first is that differences in upward occupational mobility are strongly persistent over time. While at first glance, this might not seem particularly striking given only a 20-year gap, we find this surprising given that there were significant changes over the period. The industrial make-up of the country changed dramatically, with some areas being hit particularly hard by the decline of manufacturing, mining, and fishing. While some areas saw industrial decline, the "big-bang" in the 1980s led to London becoming a major global financial centre. Despite this, in general, the upward mobility hotspots for those born in the late 1970s were the same as for those born in the late 1950s.

Second, measures of absolute and relative mobility tend to be spatially positively correlated, subject to the discussion on the extent to which drawing such a conclusion emerges for different mobility metrics. Indeed, the relationship is strongest and interpretation clearest for homeownership. In terms of probability of purchasing a home, area matters significantly more for children of non-homeowner parents. Previous studies have shown the equivalent result for income, and the evidence here suggests that heterogeneous area effects could be even more important for the understanding of wealth mobility.

Third, there is a strong relationship between upward education and upward occupation mobility. While in the current analysis it is not possible to show this is causal, it at least suggests that proxies for occupation or wage mobility based on education measures might be valid. Finally, the finding that upward homeownership mobility and upward occupation mobility are weakly negatively correlated is novel and potentially important for the way researchers and policymakers frame geographical differences in mobility. Areas that score highly on an occupation or wage-based measure of mobility might score poorly on homeownership mobility, which can be considered a proxy for wealth mobility. Unless one is willing to make value judgements on the relative importance of various dimensions of mobility, it will generally be difficult to rank areas in terms of overall mobility.

To conclude, the focus of the paper has been to generate consistent and robust estimates of geographical differences in intergenerational mobility for different areas of England and Wales with the best data available. A particular emphasis is placed on the comparison of London and elsewhere. Inspired by the recent US experience, in which initial descriptive work led to ample secondary analysis to inspect for causality, this offers a first step towards a greater understanding of intergenerational mobility in the UK. In terms of future work, one promising direction would be to use the regional changes in mobility rates over time and to consider reforms that affected areas to different degrees. The key challenge here will be working with a small number of subnational areas, and there would be significant value in developing new datasets, such as matched tax records, which could deliver finer geographical estimates. In addition, the analysis of the paper does not differentiate between "movers" and "stayers". In areas with low upward mobility, do those individuals who stay as adults have the same poor outcomes as those who move away? As a future line of investigation, this would be valuable for enhancing understanding of why the place where people grow up seems to matter so much for later outcomes. Finally, we see our result that there is a weak negative correlation between upward homeownership and occupational mobility as worthy of further empirical and theoretical work on the interaction between wealth and earnings/occupational mobility.

Supporting information

Additional supporting information can be found online in the supporting information section at the end of the article.

Online appendix
Replication files

References

Acciari, P., Polo, A., and Violante, G. (2019), 'And yet, it moves': intergenerational mobility in Italy, NBER Working Paper 25732.

Alexander, J. T., Gardner, T., Massey, C., and O'Hara, A. (2015), Creating a longitudinal data infrastructure at the Census Bureau, Working Paper, Center for Administrative Records Research and Applications, US Census Bureau.

Black, S. and Devereux, P. (2011), Recent developments in intergenerational mobility, in D. Card and O. Ashenfelter (eds), *Handbook of Labor Economics*, Volume 4B, North-Holland, Amsterdam.

Blackwell, L., Lynch, K., Smith, J., and Goldblatt, P. (2003), Longitudinal study 1971–2001: completeness of census linkage, Series LS no. 10, Office for National Statistics.

Blanden, J. (2013), Cross-country rankings in intergenerational mobility: a comparison of approaches from economics and sociology, *Journal of Economic Surveys 27*, 38–73.

Blanden, J. and Machin, S. (2004), Educational inequality and the expansion of UK higher education, *Scottish Journal of Political Economy 51*, 230–249.

Blanden, J. and Machin, S. (2008), Up and down the generational income ladder in Britain: past changes and future prospects, *National Institute Economic Review 205*, 101–116.

Blanden, J. and Machin, S. (2017), Home ownership and social mobility, Centre for Economic Performance Discussion Paper 1466.

Blanden, J., Goodman, A., Gregg, P., and Machin, S. (2004), Changes in intergenerational mobility in Britain, in M. Corak (ed.), *Generational Income Mobility*, Cambridge University Press, Cambridge.

Blanden, J., Gregg, P., and Machin, S. (2005), Intergenerational mobility in Europe and North America, Report supported by the Sutton Trust, Centre for Economic Performance, London School of Economics.

Blanden, J., Gregg, P., and Macmillan, L. (2013), Intergenerational persistence in income and social class: the effect of within-group inequality, *Journal of the Royal Statistical Society: Series A (Statistics in Society) 176*, 541–563.

Blanden, J., Eyles, A., and Machin, S. (2021), Trends in intergenerational homeownership and wealth transmission, Centre for Economic Performance Discussion Paper 1756.

Boserup, S., Kopczuk, W., and Kreiner, C. (2013), Intergenerational wealth mobility: evidence from Danish wealth records of three generations, University of Copenhagen mimeo.

Bukodi, E. and Goldthorpe, J. (2018), *Social Mobility and Education in Britain: Research, Politics and Policy*, Cambridge University Press, Cambridge.

Buscha, F. and Sturgis, P. (2018), Declining social mobility? Evidence from five linked censuses in England and Wales 1971–2011, *British Journal of Sociology 69*, 154–182.

Butikofer, A., Dalla Zuanna, A., and Salvanes, K. (2018), Breaking the links: natural resource booms and intergenerational mobility, NHH Department of Economics Discussion Paper 19.

Card, D. and Lemieux, T. (2001), Can falling supply explain the rising return to college for younger men? A cohort-based analysis, *Quarterly Journal of Economics 116*, 705–746.

Card, D., Domnisoru, C., and Taylor, L. (2018), The intergenerational transmission of human capital: evidence from the golden age of upward mobility, NBER Working Paper 25000.

Chetty, R. and Hendren, N. (2018a), The impacts of neighborhoods on intergenerational mobility I: childhood exposure effects, *Quarterly Journal of Economics 133*, 1107–1162.

Chetty, R. and Hendren, N. (2018b), The impacts of neighborhoods on intergenerational mobility I: county-level estimates, *Quarterly Journal of Economics 133*, 1163–1228.

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014), Where is the land of opportunity? The geography of intergenerational mobility in the United States, *Quarterly Journal of Economics 129*, 1553–1623.

Connelly, R., Gayle, V., and Lambert, P. (2016), A review of occupation-based social classifications for social survey research, *Methodological Innovations 9*, 1–14.

Corak, M. (2020), The Canadian geography of intergenerational income mobility, *Economic Journal 130*, 2134–2174.

Corak, M. and Heisz, A. (1999), The intergenerational earnings and income mobility of Canadian men: evidence from longitudinal income tax data, *Journal of Human Resources 34*, 504–533.

Dahl, M. and DeLeire, T. (2008), The association between children's earnings and fathers' lifetime earnings: estimates using administrative data, University of Wisconsin-Madison, Institute for Research on Poverty.

Dearden, L., Machin, S., and Reed, H. (1997), Intergenerational mobility in Britain, *Economic Journal 107*, 47–66.

Deutscher, N. and Mazumder, B. (2019), Intergenerational mobility in Australia: national and regional estimates using administrative data, Life Course Centre Working Paper 2019-02.

Deutscher, N. and Mazumder, B. (2021), Measuring intergenerational income mobility: a synthesis of approaches, *Journal of Economic Literature*, forthcoming.

Elliot Major, L. and Machin, S. (2018), *Social Mobility and Its Enemies*, Pelican Books, London.

Erikson R. and Goldthorpe, J. (2010), Has social mobility in Britain decreased? Reconciling divergent findings on income and class mobility, *British Journal of Sociology 61*, 211–230.

Eurostat (2019), GDP at current market prices by NUTS 2 regions, series nama_10r_2gdp, http://appsso.eurostat.ec.europa.eu/.

Friedman, S. and Macmillan, L. (2017), Is London really the engine-room? Migration, opportunity hoarding and regional social mobility in the UK, *National Institute Economic Review 240*, 58–72.

Heidrich, S. (2017), Intergenerational mobility in Sweden: a regional perspective, *Journal of Population Economics 30*, 1241–1280.

Jäntti, M. and Jenkins, S. (2013), Income mobility?, IZA Discussion Paper 7730.

Jonsson, J., Grusky, D., Di Carlo, M., Pollak, R., and Brinton, M. (2009), Microclass mobility: social reproduction in four countries, *American Journal of Sociology 114*, 977–1036.

Lambert, P. and Bihagen, E. (2014), Using occupation-based social classifications, *Work, Employment and Society 28*, 481–494.

Lambert, P. and Griffiths, D. (2018), *Social Inequalities and Occupational Stratification: Methods and Concepts in the Analysis of Social Distance*, Palgrave, Basingstoke.

Lambert, P., Prandy, K., and Bottero, W. (2007), By slow degrees: two centuries of social reproduction and mobility in Britain, *Sociological Research Online 12*, 1–26.

Mazumder, B. and Acosta, M. (2015), Using occupation to measure intergenerational mobility, *Annals of the American Academy of Political and Social Science 657*, 174–193.

Platt, L. (2005), The intergenerational social mobility of minority ethnic groups, *Sociology 39*, 445–461.

Platt, L. (2007), Making education count: the effects of ethnicity and qualifications on intergenerational social class mobility, *Sociological Review 55*, 485–508.

Rohenkohl, B. (2019), Intergenerational income mobility in the UK, Sheffield Economic Research Paper 2019017.

Savage, M. et al. (2013), A new model of social class? Findings from the BBC's Great British Class Survey experiment, *Sociology 47*, 219–250.

Stewart, A., Prandy, K., and Blackburn, R. (1980), *Social Stratification and Occupations*, MacMillan, London.

Sobek, M. (1995), The comparability of occupations and the generation of income scores, *Historical Methods: A Journal of Quantitative and Interdisciplinary History 28*, 47–51.

Social Mobility Commission (2017), State of the Nation 2017: Social Mobility in Great Britain, UK Government report, https://www.gov.uk/government/publications/state-of-the-nation-2017.

Tan, H. (2018), A different land of opportunity: the geography of intergenerational mobility in the early 20th-century US, mimeo.

First version submitted December 2020;
final version received July 2022.