

Small Area Consumption Estimates Combining Survey and Financial Footprints Data

Peter Levell, Lars Nesheim and Gautam Vyas

ESCoE Discussion Paper 2025-03 (Revised)

April 2026

Original version published April 2025

ISSN 2515-4664

DISCUSSION PAPER

Abstract

We use small-area estimation methods that combine information from a household budget survey, a much larger survey of local demographics and employment, and area-level information on bank account outflows and energy consumption to estimate average equivalised consumption measures across 367 local authority districts in Great Britain. We show that including bank account data substantially improves our estimates, showing that these and other financial footprints data can play an important role in measuring local consumption and hence living standards. We also compare consumption measures that correspond to welfare under different assumptions about mobility and the capitalisation of local amenities into house prices, as well as traditional local income measures, and show that the rankings of local authorities are sensitive to the choice of measure.

Keywords: Small area estimation, imputation, local consumption

JEL classification: E21, I32, R12

Peter Levell, IFS
peter.l@ifs.org.uk

Published by:
Economic Statistics Centre of Excellence
King's College London
Strand
London
WC2R 2LS
United Kingdom
www.escoe.ac.uk

ESCoE Discussion Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Economic Statistics Centre of Excellence (ESCoE), its partner institutions or the Office for National Statistics (ONS).

Small Area Consumption Estimates Combining Survey and Financial Footprints Data

Peter Levell^{*}, Lars Nesheim[†] and Gautam Vyas[‡]

April, 2026

Abstract

We use small-area estimation methods that combine information from a household budget survey, a much larger survey of local demographics and employment, and area-level information on bank account outflows and energy consumption to estimate average equivalised consumption measures across 367 local authority districts in Great Britain. We show that including bank account data substantially improves our estimates, showing that these and other financial footprints data can play an important role in measuring local consumption and hence living standards. We also compare consumption measures that correspond to welfare under different assumptions about mobility and the capitalisation of local amenities into house prices, as well as traditional local income measures, and show that the rankings of local authorities are sensitive to the choice of measure.

Keywords: Small area estimation, imputation, local consumption

JEL classification: E21, I32, R12

Acknowledgements: This paper was produced as part of the project ‘Towards a single index of local and national sustainable wellbeing’ funded by the Economic Statistics Centre of Excellence. Funding the ESRC-funded Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (grant reference ES/Z504634/1) is also gratefully acknowledged. All errors and omissions remain the responsibility of the authors. This paper makes use of data on bank account outflows kindly shared with the IFS from the Financial Data Service (FINDS). FINDS bears no responsibility for any errors in the interpretation of their data and the views expressed in this paper are our own. This work was undertaken in the Office for National Statistics Secure Research Service using data from ONS and other owners and does not imply the endorsement of the ONS or other data owners. Correspondence: peter.l@ifs.org.uk, gautam.vyas@ifs.org.uk

^{*}Institute for Fiscal Studies and ESCoE

[†]Institute for Fiscal Studies, UCL, University of Bristol, and ESCoE

[‡]Institute for Fiscal Studies and ESCoE

1 Introduction

Spatial inequalities are a key public concern in many countries. For example, surveys of public attitudes rank inequalities across areas as the most significant form of inequality in the UK: ahead of inequalities by income, race, gender or across generations (Benson et al., 2024). This means there is enormous policy interest in measuring and tracking these inequalities over time. Area-level inequalities are typically measured in terms of area differences in incomes or output (for example McCann (2020) compares regional inequality across countries using 28 measures based on regional disposable income or output per head). But measures of output reflect the productivity of those working in a location, rather than residents, and current incomes may not be accurate measures of households’ resources if individuals borrow or dissave.

Economists have long argued that there are strong theoretical reasons to view current consumption as a better measure of households’ lifetime resources than current income (Poterba, 1989; Slesnick, 1993). However, while census and administrative data are increasingly used to provide fine-grained information on local incomes, much less information is available to produce reliable estimates of local consumption. Sample sizes in household budget surveys are typically too small to generate precise local-area averages. At the same time, various forms of naturally occurring data that track households’ ‘financial footprints’ – such as card transactions or current account outflows – contain valuable information on spending, but cannot on their own be used to construct measures consistent with national accounts definitions of consumption, which include imputed service flows (for example, for owner-occupied housing). Moreover, these data are not typically linked to the household characteristics required to adjust (‘equivalise’) spending for differences in family size and composition.

In this paper, we show how financial footprints data can be combined with small-area estimation techniques to produce reliable estimates of local average consumption spending. Our contribution is twofold. Firstly, we quantify the scale of improvement from incorporating financial footprints data into these estimates. Secondly, we compare empirical estimates of local living standards estimated using different definitions of consumption – each corresponding to economic welfare under different assumptions about the amenity capitalisation into house prices – with each other and with more traditional measures of average individual incomes.

To do this, we produce estimates of resident households’ average consumption spending for 367 local authority districts across Great Britain by combining information from a household budget survey with an annual sample size of 5,000 households

with information on local demographic characteristics from a much larger population survey, area-level information on average outflows from current accounts, and household energy use.

We use simulations to evaluate different methods and show that ‘Empirical Best’ (EB) small-area estimation techniques ([Molina and Rao, 2010](#); [Pfeffermann, 2013](#)) greatly improve performance relative to direct measures of the mean using the budget survey alone, as well as naive regression imputation methods that rely solely on consumption proxies. The average improvement from using financial footprints data varies according to the definition of consumption we use, and is greatest when we exclude housing costs or deflate housing expenditures according to local house prices (when it averages 9% across all areas). This suggests that such data can be usefully combined with survey data to improve the quality of statistical estimates.

The case for estimating household consumption rests on it being a better measure of economic well-being than income. One recent article by [Meyer and Sullivan \(2023\)](#) puts this case well: “Consumption better reflects long-run resources and is more likely to capture disparities that result from differences across families in the accumulation of assets or access to credit. Consumption will reflect the loss of housing service flows if home ownership falls, the loss in wealth if asset values fall, and the belt-tightening that a growing debt burden might require, all of which an income measure would miss.” These considerations mean that rankings of households by standard of living can differ significantly depending on whether they are calculated using consumption or income ([Blundell and Preston, 1995](#)). For example, retired households are likely to have lower than average incomes, but can often maintain high rates of consumption by drawing down their wealth. This will translate into differences in area-level measures of living standards across areas with younger or older populations (e.g. student towns compared to popular retirement locations). Areas that see significant wealth gains - because of rising local house prices for example - may also see increases in their relative living standards that will not necessarily be reflected in average local incomes.¹

Despite these advantages, simple comparisons of average nominal expenditure across local areas may not perfectly capture economic well-being ([Diamond and Moretti, 2021](#)). This is because local prices and non-market amenities will differ across areas. Consumption measures that adjust for local prices or housing costs might therefore better capture material living standards, though which measure is

¹Another reason to favour consumption spending is that it may be less affected by reporting biases and measurement errors (particularly at the bottom of the income distribution, [Brewer et al. \(2017\)](#)).

best depends on assumptions about utility functional form, local amenity and house price variation, and mobility frictions.

For example, if unobserved non-market amenities vary across local areas and affect the marginal utility of consumption, spatial differences in consumption baskets will, in part, reflect the effects of amenities rather than directly reflecting welfare differences. Furthermore, these unobserved amenities will influence the price of housing (and possibly of other services) with the degree of capitalisation depending on the degree of household mobility. Nominal consumption expenditure measurements that include housing-related expenditures are equivalent to welfare under the assumptions of the Rosen-Roback model ([Rosen, 1979](#); [Roback, 1982](#)) (homogenous preferences and perfect mobility). Under these assumptions, amenity-adjusted housing prices are equal across areas, with observed differences in the price of housing (and other consumption goods) exactly reflecting differences in the value of local amenities. In contrast, consumption that has been deflated according to local house prices will reflect welfare under the assumption that amenity flows are constant across areas and prohibitively high mobility frictions (or idiosyncratic preferences for location) prevent prices from equalising across areas. Finally, nominal consumption excluding housing-related expenditures are equivalent to welfare if housing service flows and the price of non-housing goods are equal across areas.

Since each of these consumption measures has strengths and limitations as a proxy for welfare, we report results for all of them and discuss how they differ. We find that the inclusion of nominal housing-related expenditures increases the ranking of several local authorities in London. However, deflating consumption results in large relative falls for these areas, indicating that high nominal consumption in these areas reflects high house prices. We calculate summary measures of inequality across local areas and find that inequality is greatest for nominal consumption including housing, whilst inequality in deflated consumption is less pronounced. Conclusions about geographic disparities that policymakers might draw from these results, depend on how important they believe preference heterogeneity and mobility frictions are. If the assumptions of the Rosen-Roback model hold, high prices and nominal expenditures in London suggest individuals in these areas enjoy higher welfare from amenity flows than their counterparts in other areas of the country. In contrast, if amenity flows are similar across areas, the presence of mobility frictions may prevent individuals in high house-price areas (such as London) from moving elsewhere to obtain higher welfare.

We also compare our measures of average equivalised consumption with per capita income measures that have also been used to compare local living standards ([Judge and McCurdy \(2022\)](#)). We show that the rank of local authorities can differ

substantially across these two measures, with several areas of London ranking much higher in terms of per capita income than they do on equivalised average deflated consumption including housing related expenditures. Local authorities in London have incomes on average 40% above the Great Britain mean, yet their deflated consumption including housing related expenditures is 2.97% below the national mean.

The rest of this paper is structured as follows. In Section 2 we describe small area estimation methods and the particular approaches we will use to estimate area-level average consumption. Section 3 describes the data we use. Section 4 presents our empirical results providing a detailed analysis of geographic disparity of consumption and income across UK local authorities. Section 5 concludes.

2 Small area estimation

Our goal is to estimate the mean of equivalised consumption in all areas $a \in \{1, \dots, A\}$:

$$\bar{c}_a = \mathbb{E}[c_{a,i}] = \mathbb{E}\left[\frac{C_{a,i}}{\Omega(S_i)}\right],$$

where the expectation is with respect to individual households indexed by i , $c_{a,i}$ is equivalised household consumption, $C_{a,i}$ is unadjusted household consumption, and $\Omega(S_i)$ is an equivalisation scale based on S_i , the family size of household i .

Denote the population size of area a as N_a .² Let the N_a -dimensional vector of equivalised consumption in area a be $\mathbf{c}_a = (\mathbf{c}'_{a,r}, \mathbf{c}'_{a,s})$, where $\mathbf{c}_{a,s}$ is an n_a -dimensional vector of equivalised consumption observed in a survey dataset, and $\mathbf{c}_{a,r}$ is the remaining vector of unobserved equivalised consumption values for households who did not participate in the survey. Define the vectors of unequivalised consumption analogously such that $\mathbf{C}_a = (\mathbf{C}'_{a,r}, \mathbf{C}'_{a,s})$. Let I_r and I_s be the sets of indices in the missing and observed subvectors respectively.

The most straightforward estimator for $\bar{c}_a$ is the simple survey sample average also called the *direct* estimator calculated using the survey data alone. This estimator is:

$$\hat{\bar{c}}_a^{Direct} = \frac{1}{n_a} \sum_{i \in I_s}^{n_a} c_{a,i}. \quad (2.1)$$

²In typical applications, N_a is the true population size, and demographic information on all individuals is available in a census dataset. Because our aim is to estimate average consumption for local authorities in 2018/2019 (and there was no census in this year), we instead use the combined samples of households from a large survey (the Annual Population Survey) and a smaller budget survey as our ‘population’ dataset (with full details in Section 3), so N_a will refer to the total number of households in area a in both of these surveys.

While this estimator is unbiased for the population mean (assuming random sampling within areas), it is likely to be very noisy when sample sizes are small. Moreover, it is not possible to estimate this quantity in areas for which the sample size is zero.

To overcome these limitations, small area estimation methods bring in additional data from auxiliary sources to complement the information in smaller surveys (Pfeffermann, 2013). This additional data could be a vector of covariates $\mathbf{X}_a = (\mathbf{X}'_{a,r}, \mathbf{X}'_{a,s})$ for individuals in each area from a population census or another survey (in which consumption is not observed), and/or area-level variables from some other dataset, $\mathbf{Z}_a$. In particular, we calculate

$$\hat{c}_a^{EB} = \frac{n_a}{N_a} \hat{c}_a^{Direct} + \frac{1}{N_a} \sum_{i \in I_r} \mathbb{E}[c_{a,i} | \mathbf{c}_{a,s}, \mathbf{X}_a, \mathbf{Z}_a]. \quad (2.2)$$

This small area estimator improves upon the direct estimator when n_a is small by combining it with a large sample regression estimator. To make use of this small area estimator, we need to estimate a model for the conditional expectation $\mathbb{E}[c_{a,i} | \mathbf{c}_{a,s}, \mathbf{X}_a, \mathbf{Z}_a]$. For this we use the nested error linear regression model (Battese et al., 1988).

2.1 Nested error linear regression model for log consumption

Assume the following linear regression model for log consumption:

$$\begin{aligned} \log \mathbf{C}_{a,r} &= \mathbf{X}_{a,r} \boldsymbol{\beta} + \mathbf{Z}_a \boldsymbol{\pi} + u_a \mathbf{1}_{N_a - n_a} + \boldsymbol{\varepsilon}_{a,r} \\ \log \mathbf{C}_{a,s} &= \mathbf{X}_{a,s} \boldsymbol{\beta} + \mathbf{Z}_a \boldsymbol{\pi} + u_a \mathbf{1}_{n_a} + \boldsymbol{\varepsilon}_{a,s} \\ u_a &\sim_{iid} N(0, \sigma_u^2) \\ \varepsilon_{i,a} &\sim_{iid} N(0, \sigma_\varepsilon^2). \end{aligned} \quad (2.3)$$

This model includes individual level covariates $(X_{a,r}, X_{a,s})$, area level covariates Z_a , area-specific random effects u_a , and individual level errors $\varepsilon_{i,a}$.³

The regression model implies:

$$\begin{pmatrix} \log \mathbf{C}_{a,r} \\ \log \mathbf{C}_{a,s} \end{pmatrix} \sim N \left[\begin{pmatrix} \mathbf{X}_{a,r} \boldsymbol{\beta} + \mathbf{Z}_a \boldsymbol{\pi} \\ \mathbf{X}_{a,s} \boldsymbol{\beta} + \mathbf{Z}_a \boldsymbol{\pi} \end{pmatrix}, \begin{pmatrix} \sigma_u^2 \mathbf{1}_{N_a - n_a} \mathbf{1}'_{N_a - n_a} + \sigma_\varepsilon^2 \mathbf{I}_{N_a - n_a} & \sigma_u^2 \mathbf{1}_{N_a - n_a} \mathbf{1}'_{n_a} \\ \sigma_u^2 \mathbf{1}_{n_a} \mathbf{1}'_{N_a - n_a} & \sigma_u^2 \mathbf{1}_{n_a} \mathbf{1}'_{n_a} + \sigma_\varepsilon^2 \mathbf{I}_{n_a} \end{pmatrix} \right].$$

³Small area estimation methods are usually divided into those that use individual level covariates (e.g. Battese et al., 1988) and those that use area level information (e.g. Fay and Herriot, 1979). In our approach, we use both types of covariates.

which further implies $\log \mathbf{C}_{a,r} | \log \mathbf{C}_{a,s} \sim N(\boldsymbol{\mu}_{a,r|s}, \mathbf{V}_{a,r|s})$, where

$$\begin{aligned}\boldsymbol{\mu}_{a,r|s} &= \mathbf{X}_{a,r}\boldsymbol{\beta} + \mathbf{Z}_a\boldsymbol{\pi} + \sigma_u^2 \mathbf{1}_{N_a-n_a} \mathbf{1}'_{n_a} (\sigma_u^2 \mathbf{1}_{n_a} \mathbf{1}'_{n_a} + \sigma_\varepsilon^2 \mathbf{I}_{n_a})^{-1} (\log \mathbf{C}_{a,s} - \mathbf{X}_{a,s}\boldsymbol{\beta} - \mathbf{Z}_a\boldsymbol{\pi}) \\ &= \mathbf{X}_{a,r}\boldsymbol{\beta} + \mathbf{Z}_a\boldsymbol{\pi} + \frac{\gamma_a}{n_a} \mathbf{1}_{N_a-n_a} \mathbf{1}'_{n_a} (\log \mathbf{C}_{a,s} - \mathbf{X}_{a,s}\boldsymbol{\beta} - \mathbf{Z}_a\boldsymbol{\pi}) \\ \mathbf{V}_{a,r|s} &= \sigma_u^2 (1 - \gamma_a) \mathbf{1}_{N_a-n_a} \mathbf{1}'_{N_a-n_a} + \sigma_\varepsilon^2 \mathbf{I}_{N_a-n_a},\end{aligned}$$

with $\gamma_a = \sigma_u^2 (\sigma_u^2 + \sigma_\varepsilon^2/n_a)^{-1}$.

For each element of $\log \mathbf{C}_{a,r}$, the conditional mean is the sum of individual effects, area level effects, and the average residuals in area a . The impacts of the average model residuals are shrunk towards zero when the local sample size is small or when the idiosyncratic errors account for a large share of the residual variance.

These expressions allow us to obtain the conditional mean of $\log$ consumption for each area. Jensen's inequality implies that the exponential of this less than the mean of consumption. To obtain values for mean consumption, we follow [Molina and Rao \(2010\)](#) and use a Monte-Carlo approximation. This involves drawing L vectors of $\log \mathbf{C}_{a,r}$ for each area from the normal distribution with area-specific conditional mean and variance, exponentiating, and then averaging.⁴ Doing this directly would involve simulation of A vectors each with $N_a - n_a$ elements. When areas have large populations, repeated simulation of all the non-sampled units across many bootstrap replications would be computationally burdensome. Instead of drawing from multivariate distributions, they suggest using draws from the model

$$\log \mathbf{C}_{a,r} = \boldsymbol{\mu}_{a,r|s} + v_a \mathbf{1}_{N_a-n_a} + \boldsymbol{\varepsilon}_{a,r},$$

where $v_a \sim N(0, \sigma_u^2 (1 - \gamma_a))$ and $\boldsymbol{\varepsilon}_{a,r} \sim N(\mathbf{0}_{N_a-n_a}, \sigma_\varepsilon^2 \mathbf{I}_{N_a-n_a})$. This has the same distribution, but has the computational advantage of only requiring draws from independent univariate distributions.

In summary, for $l = 1, \dots, L$, we draw $\widehat{\log C_{a,i}}^\ell$ from the above distribution, exponentiate to produce $\hat{C}_{a,i}^\ell = \exp\left(\widehat{\log C_{a,i}}^\ell\right)$, and then compute the average in each area:

$$\hat{\bar{C}}_a^{EB} = \frac{1}{L} \sum_{\ell=1}^L \left(\frac{1}{N_a} \left[\sum_{i \in I_r} \frac{\hat{C}_{a,i}^\ell}{\Omega(S_i)} + \sum_{i \in I_s} \frac{C_{a,i}}{\Omega(S_i)} \right] \right),$$

We refer to this procedure as the Empirical Best (EB) estimator, because, under model (2.3), it minimises mean-squared error over the class of all possible estimators $\hat{\bar{C}}_a$ ([Molina and Rao, 2010](#)).

⁴Since $C_{a,r}$ is log-normal, an alternative method to obtain the conditional mean of consumption would be to use the identity: $\mathbb{E}(\mathbf{C}_{a,r} | \mathbf{C}_{a,s}, \mathbf{X}_{a,r}, \mathbf{Z}_a) = e^{\boldsymbol{\mu}_{a,r|s} + 0.5 \times \text{diag}(\mathbf{V}_{a,r|s})}$.

In areas where $n_a = 0$, the EB estimator reduces to a regression prediction based only on $\mathbf{X}_{a,r}$, $\mathbf{Z}_a$ and draws of the household specific residual $\boldsymbol{\varepsilon}_a$. This “pure” regression prediction is commonly referred to as a *synthetic* estimator; we henceforth refer to it as a linear regression estimator. In our empirical application, we also compute estimates of mean consumption using the linear regression estimator and compare them to the EB estimates.

2.2 MSE estimation

We use two approaches to estimate the MSE’s associated with the EB and alternative estimation strategies. The first is a model-based simulation approach. The second is a design-based approach based on repeated draws from a fixed population.

2.2.1 Model-based MSEs

The model based MSE of $\hat{c}_a^{EB}$ is

$$MSE_{\mathcal{M}}(\hat{c}_a^{EB}) = \mathbb{E}_{\mathcal{M}} \left(\hat{c}_a^{EB} - \bar{c}_a \right)^2, \quad (2.4)$$

where the expectation $\mathbb{E}_{\mathcal{M}}$ is taken over realisations of the random effects and errors $(\mathbf{u}_a, \boldsymbol{\varepsilon}_a)$. Model based MSE’s for the other estimators are analogous. Under the model based approach, both the finite population and the realised sample are treated as random objects generated by the hierarchical model in (2.3).

We estimate the model based MSE’s using a parametric bootstrap following the steps outlined in [Molina and Rao \(2010\)](#) (who follow the bootstrap method of [González-Manteiga et al., 2008](#)). The steps of the procedure are:

1. Estimate model (2.3) by applying maximum likelihood estimation to the survey sample. Obtain estimates of $\boldsymbol{\beta}, \boldsymbol{\pi}, \sigma_u^2, \sigma_\varepsilon^2$, denoted: $\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\pi}}, \hat{\sigma}_u, \hat{\sigma}_\varepsilon$.
2. Using these estimates, draw B i.i.d bootstrapped populations with size N_a in each area $\{\log c_{i,a}^{*(b)}; i = 1, \dots, N_a, a = 1, \dots, A\}$ from the following super-population model

$$\log \mathbf{C}_a^* = \mathbf{X}_a \hat{\boldsymbol{\beta}} + \mathbf{Z}_a \hat{\boldsymbol{\pi}} + u_a^* \mathbf{1}_{N_a} + \boldsymbol{\varepsilon}_a^*; \quad u_a^* \sim N(0, \hat{\sigma}_u^2), \boldsymbol{\varepsilon}_{i,a}^* \sim N(0, \hat{\sigma}_\varepsilon^2).$$

Taking the covariates for the first $N_a - n_a$ observations from the census data and the final n_a observations from the smaller survey dataset.

3. For each bootstrapped population (indexed b), calculate the area-level averages

$$\bar{c}_a^{*(b)} = \frac{1}{N_a} \sum_{i=1}^{N_a} \exp(\log C_{a,i}^{*(b)}) / \Omega(X_i).$$

4. For each bootstrapped population, take the bootstrapped consumption observations from the households in the final n_a elements (that is, the observations from the smaller survey sample), and use these to produce the EB estimators as described in Section 2.1, linear regression estimators and direct estimators. For comparison, also, estimate a version of the EB estimator that omits data on $\mathbf{Z}_a$. Denote the resulting estimates as $\hat{c}_a^{EB*(b)}$, $\hat{c}_a^{LR*(b)}$, $\hat{c}_a^{Direct*(b)}$ and $\hat{c}_a^{noZ*(b)}$.
5. Calculate the estimate of the mean squared error across the bootstrap draws for each area. For example, for the EB estimator, this is

$$MSE_{\mathcal{M}}^*(\hat{c}^{EB}) = \frac{1}{B} \sum_{b=1}^B (\hat{c}_a^{EB*(b)} - \bar{c}_a^{*(b)})^2.$$

The approximation is analogous for the other estimators.

The estimator that omits information on $\mathbf{Z}_a$, by necessity, performs no better than EB. We use the comparison to show empirically the degree to which financial footprints and energy use data improve the accuracy of our small-area estimates.

2.2.2 Design-based MSEs

An alternative design-based MSE aims to measure

$$MSE_{\mathcal{D}}(\hat{c}_a^{EB}) = \mathbb{E}_{\mathcal{D}} \left[(\hat{c}_a^{EB} - \bar{c}_a)^2 | \mathcal{P} \right]$$

where the expectation is taken over alternative draws of the sample I_s (and conditional on the fixed finite population $\mathcal{P}$). Under the design-based MSE definition, the finite population is treated as fixed and only the sampling indicators are random.

We approximate the design-based MSE using random subsampling across areas with sample sizes greater than 20 in our consumption survey. Specifically, for draws $d = 1 \dots D$, we do the following:

1. For each area a , randomly partition the observations in our small consumption survey sample into an estimation sample and a holdout sample, each of size $\frac{n_a^*}{2}$ (where $n_a^* = n_a - 1$ if n_a is odd, and $n_a^* = n_a$ otherwise). Index observations in the estimation sample in draw d by $I_s^{(d)}$ and denote the vector of consumption values in this sample by $\mathbf{C}_{a,s^{(d)}}$.

2. Then calculate direct, EB and linear regression estimates of average consumption for each area and calculate the differences between these estimates and the means from the holdout sample. For the EB method, estimate

$$\log \mathbf{C}_{a,s(d)} = \mathbf{X}_{a,s(d)}\boldsymbol{\beta} + \mathbf{Z}_a\boldsymbol{\pi} + u_a \mathbf{1}_{\frac{n_a^*}{2}} + \boldsymbol{\varepsilon}_{a,s(d)},$$

using observations $i \in I_s^{(d)}$, and use the results to impute $\mathbf{C}_{a,r}$ for the $N_a - n_a$ observations with missing consumption information. Then use in combination with the observations in $I_s^{(d)}$ to compute $\hat{c}_a^{EB(d)}$. Also compute the direct estimator

$$\hat{c}_a^{Direct(d)} = \frac{1}{2n_a^*} \sum_{i \in I_s^{(d)}} c_i.$$

3. Draw a new estimation sample in each area and repeat.

The average squared deviations from this procedure yields our design-based estimate of the MSE. For the EB estimator, the MSE estimate is:

$$\widehat{MSE}_{\mathcal{D}}(\hat{c}_a^{EB}) = \frac{1}{D} \sum_{d=1}^D \left(\hat{c}_a^{EB(d)} - \frac{1}{n_a^*/2} \sum_{i \notin I_s^d} c_i \right)^2.$$

2.3 Standard errors

The model-based MSE's above incorporate both sampling uncertainty and uncertainty in the values of the hyper-parameters $\mathbf{u}_a$. To calculate standard errors for our estimates that remove the variability associated with different draws of the hyper-parameters (i.e. apply for a single population), we adopt a different bootstrap procedure. For the EB estimator this is:

1. Estimate model (2.3) applying maximum likelihood to the survey sample. Obtain estimates of $\boldsymbol{\beta}, \boldsymbol{\pi}, \mathbf{u}_a, \sigma_\varepsilon^2$: $\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\pi}}, \hat{\mathbf{u}}_a$ and $\hat{\sigma}_\varepsilon$.
2. Take draws of from the distribution of individual-level residuals $\varepsilon_{i,a}^{\dagger(b)} \sim N(0, \hat{\sigma}_\varepsilon^2)$.
3. Use these draws to simulate $\log C_{a,i}^{\dagger(b)}$ using:

$$\log C_{a,i}^{\dagger(b)} = \mathbf{X}_{a,i}\hat{\boldsymbol{\beta}} + \mathbf{Z}_a\hat{\boldsymbol{\pi}} + \hat{u}_a + \varepsilon_{i,d}^{\dagger(b)}. \quad (2.5)$$

4. For each bootstrapped population, take observations in I_s , use these to compute EB estimates as described in Section 2.1, $\hat{c}_a^{EB\dagger(b)}$.

The standard error associated with each area level estimate is then

$$SE(\hat{c}_a^{EB}) = \sqrt{\frac{1}{B} \sum_{b=1}^B \left(\hat{c}_a^{EB\uparrow(b)} - \hat{c}_a^{EB} \right)^2}.$$

The standard errors for the other estimators are computed in an analogous way (replacing the model in (2.5) with the appropriate estimation model).

3 Data and consumption measurement

Living Costs and Food Survey (LCFS): Our source of consumption data is the LCFS. The LCFS is an annual survey recording the spending patterns of a sample of households in Great Britain ([Office for National Statistics, 2021](#)). We pool data for the calendar years 2018 and 2019, which we obtain from the 2017, 2018 and 2019 datasets (which cover the financial years 2017-2018, 2018-19 and 2019-20 respectively). Households in the survey record small expenditures in a spending diary over the course of two weeks, and answer recall questions on larger ‘big ticket’ expenditures such as cars and holidays. We focus on households where the household reference person (HRP) is older than 18, and remove one household with negative expenditure in the sample, as well as trimming the household expenditure distribution at the 99th percentile. We also drop a small number of households with missing covariate values. Our final sample comprises 8,744 households across the two years.

The LCFS contains aggregated COICOP⁵ category level expenditures which are derived from the spending diary and recall questions. We construct three measures of consumption from the data, described in detail in Section 3.3. Alongside consumption measures, the LCFS contains detailed information on a large set of individual and household-level demographics.

We use a secure-access, geocoded version of the LCFS with information on respondent households’ postcodes. We map these across 2019 local authority district (LAD) boundaries, and conduct our analysis using LADs as the areas described in Section 2.

Annual Population Survey: As one auxiliary data source, we use geocoded versions of the 2018 and 2019 Annual Population Survey (APS) ([Office for National Statistics, 2025](#)). The Annual Population Survey is a continuous household survey, conducted with the aim of providing information on social and labour market vari-

⁵Consumption categories are based on the COICOP (Classification of Individual Consumption According to Purpose) classification system.

ables at a local level in the years between the (decennial) UK censuses. It contains a large range of demographic information but, importantly for our purposes, does not contain information on household consumption. We apply the same sample restrictions to households and derive the same explanatory variables as in the LCFS.⁶ The final sample size after cleaning is much larger than the LCFS - around 387,000 households over the two years. As for the LCFS, we pool across two calendar years to increase the precision of our estimates.

Bank account data and energy-use: We complement the household-level covariates from the APS with area-level covariates derived from two additional auxiliary data sources. The first source is a bank account dataset provided by the Financial Data Service (FINDS). The data contains measures of average expenditure at the Middle Super Output Area (MSOA 2011) level.⁷ These measures are obtained from a sample of approximately 1.2 million bank accounts from Natwest, a large retail bank. The data covers the period May-December 2019. We use these local averages as area-level covariates in the nested error model (2.3). In addition to the APS and FINDS data, we also use data on average energy consumption at the postal outcode level⁸ published by the Department for Energy, Security and Net Zero.

Income data: We also draw upon local-authority level income data from the National Accounts Gross Disposable Household Income (GDHI) dataset, which allow us to compare the geographical distributions of consumption and income (Office for National Statistics, 2024). The GDHI is useful in our context as it is available at a granular level of geography and is largely derived from administrative data sources (Office for National Statistics, 2018). However, it is a broader measure of income than what is usually measured in surveys such as Households Below Average Income (HBAI). To convert the gross disposable income measure in the data to a cash measure (that better reflects common notions of income), we follow Judge and McCurdy (2022) in removing elements of income imputed from assets, removing deductions that would not appear on household balance sheets, and adding back mortgage interest paid by homeowners. We also compare average consumption to the 2019 Income Index of Multiple Deprivation (IMD) data, aggregated to the local

⁶Both surveys use consistent definitions of households and the household reference person. In both surveys, a “household” is “one person living alone, or a group of people (not necessarily related) living at the same address who share cooking facilities and share a living room, sitting room or dining area. A household can consist of a single family, more than one family or no families in the case of a group of unrelated people.” (Office for National Statistics, 2023b)

⁷MSOAs are UK Census defined geographic regions, smaller than local authorities.

⁸Postal outcodes are geographic areas, smaller than local authorities, corresponding to the first 2-4 digits of the UK postcode.

authority district level (available for England only, [McLennan et al. \(2019\)](#)). The Income IMD measures the share of the population in each local authority experiencing deprivation relating to low income, and is based largely on indicators of the number of families eligible for certain state benefits (such as Income Support and Jobseeker’s Allowance).

3.1 Explanatory variables

The X variables used in model (2.3) from Section 2 must be observed in both the APS and LCFS. The variables we use are: 10-year age bands of the HRP, the age at which the HRP completed their education (whether they were aged under 18, 18-21, over 21 or are still in education), an indicator for the ethnicity of the HRP, and the household’s tenure status (whether they own their homes outright, own with a mortgage or rent). We also include a year dummy for the calendar year the household was interviewed (2018 or 2019). We use information on household composition to construct dummy variables for the presence of children under 14, children aged between 14 and 18, and additional adults, as well as variables corresponding to the number of individuals from each of these groups. We also use this information to equalise our estimates of average consumption for family size, according to the OECD equivalence scale.⁹ The Z variables we use are average current account outflows/expenditures in the households’ MSOA, and average electricity consumption in KWh for the households’ local postcode.

3.2 Average consumption spending and welfare

Ultimately we are interested in calculating relative welfare in different areas. Following [Deaton and Zaidi \(2002\)](#) we adopt the measure of money metric utility, which measures the standard of living according to the level of expenditure needed to attain it. Here we extend the basic [Samuelson \(1974\)](#) framework to illustrate the role of amenity differences across geographic areas.

We start by defining the expenditure function, $e(u, p_a, \theta_a)$ as the minimum cost of attaining utility u given a set of prices p_a and a flow of local amenities θ_a , associated with area a . For clarity, we assume consumers have identical preferences. Under the assumption that the household i living in area a (with utility level $u_{a,i}$) maximises utility, consumption spending is related to prices, utility, and amenities according to

$$c_{a,i} = e(u_{a,i}, p_a, \theta_a).$$

⁹This assigns a value of 1.0 to the first adult, 0.5 to each household member aged 14 or over, and 0.3 to each child under 14.

Money metric utility uses this expression to measure utility in monetary terms by fixing the levels of prices and amenities at reference levels (p_0, θ_0) . That is, utility for household i is measured as

$$u_{a,i}^m = e(u_{a,i}, p_0, \theta_0).$$

This is simply a monotonic transformation of the original utility scale.

In terms of the money metric utility function, the utility of a person who spends $c_{a,i}$ in area a can be measured as

$$\begin{aligned} u_{a,i}^m &= e(u_{a,i}, p_a, \theta_a) \times \frac{e(u_{a,i}, p_0, \theta_0)}{e(u_{a,i}, p_a, \theta_a)} \\ &= c_{a,i} \times \frac{e(u_{a,i}, p_0, \theta_0)}{e(u_{a,i}, p_a, \theta_a)}. \end{aligned}$$

In words, welfare equals consumption spending deflated by a household specific cost of living index. The cost of living index depends on area level prices and amenities and on the specific utility level of the household under consideration.

If we assume homothetic preferences, then the expenditure function is separable and can be written $e(u_{a,i}, p_a, \theta_a) = f(u_{a,i})\tilde{e}(p_a, \theta_a)$ where f is strictly increasing. In this case, average welfare in area a can be written

$$\begin{aligned} \bar{u}_a^m &= \bar{c}_a \times \frac{\tilde{e}(p_0, \theta_0)}{\tilde{e}(p_a, \theta_a)} \\ &= \frac{\bar{c}_a}{P(p_a, \theta_a, p_0, \theta_0)} \end{aligned} \tag{3.1}$$

where $P(p_a, \theta_a, p_0, \theta_0)$ is a cost of living index that depends only on prices and amenities.

Equation (3.1) makes clear that spatial welfare comparisons based on consumption spending need to take into account of both price differences and amenity differences across areas.¹⁰ However this is not straightforward, as the monetary value consumers place on different amenities is difficult to quantify, and their relationship with local housing costs depends on assumptions about different households willingness and ability to relocate. In what follows, we therefore calculate three measures of average consumption spending which make different assumptions about $P(p_a, \theta_a, p_0, \theta_0)$.

¹⁰Here we have side-stepped a further issue of aggregation. If $f(u)$ is strictly concave, then average welfare in money metric utility is not the same as average welfare in the original utility scale since the average of $c_{a,i}$ is not the same as the average of $f(u_{a,i})$. The latter depends also on the degree of inequality in different areas.

3.3 Welfare and different consumption measures

We calculate and compare three different consumption measures.

- C1. *Non-housing consumption*: This measure includes all non-durable and durable spending but does not include spending on rent or imputed consumption flows for owner-occupiers.
- C2. *Consumption including housing*: This measure includes all spending in C1 plus spending on rent and an imputed housing consumption flow for owner-occupiers and those in social housing.¹¹ We impute housing service flows to owner occupiers using a measure of the average cost of renting houses of given types in each local authority published by the ONS. We use data from the Price Index of Private Rents (PIPR) series ([Office for National Statistics, 2019](#)). We average monthly mean rents to construct rental prices at the local authority $\times$ house type $\times$ year level and merge these into our dataset. We then add the imputed weekly housing expenditures for owners and social housing tenants and actual weekly rents for renters to the measure in C1.
- C3. *Deflated consumption spending including housing*: This measure is the same as C2 but deflates local consumption spending to account for differences in local housing costs. To deflate local consumption, we first construct local-authority level housing prices p_a^h by taking a weighted mean over house type-level prices (using national population shares in each house type from the 2021 Census ([Office for National Statistics, 2023a](#))). We then use this data to construct a Laspeyres index of the form:

$$P_a = w_h \times \frac{p_a^h}{\bar{p}^h} + (1 - w_h) \times 1$$

where w_h is the average expenditure share on housing in our data and $\bar{p}^h$ is the unweighted average across local authority level prices. In the absence of data on local non-housing prices, we assume these are invariant across areas.

Each of these measures is a valid welfare measure under a different set of assumptions. The first measure, C1, avoids making imputations, but makes the unrealistic assumption that average consumption housing service flows are the same across all local authorities.

The second measure, C2, captures welfare under the assumption that our housing service imputation model is correct and under the assumptions of the Rosen-

¹¹Those in social housing tend to pay rents below market rates and so likely enjoy a consumption flow greater than their rental expenditure implies.

Roback model. The Rosen-Roback model assumes that i) consumers have homogeneous preferences and ii) there are no costs to relocating across areas. These assumptions imply that equilibrium house prices perfectly capitalise any differences in local amenity flows and that the value of housing services and local amenity flows are captured by nominal housing spending in each location. Evidence of large mobility costs across locations in a number of contexts suggests that these assumptions also are unlikely to hold (Kennan and Walker, 2011; Bryan and Morten, 2019).

The third measure, C3, makes the assumption that amenity flows do not differ across areas. For this to coincide with persistent differences in local house prices, this requires that prohibitively high mobility costs prevent prices from equalising.

Because each of these consumption measures is an imperfect proxy for welfare, we estimate all three and compare them. In a more realistic model than the extreme assumptions implied by C2 and C3, households would face positive but not infinite mobility costs and would have heterogeneous preferences. In that case, house prices would capitalise the value of amenities as valued by the marginal mover to each area. Measuring welfare in such a model would require estimation of a full spatial equilibrium model, which is beyond the scope of the present paper.

4 Results

For each consumption measure C1 - C3, we report results for the random-effects regression models in Tables A.1 - A.3 in the Appendix. Each table includes the result from a likelihood-ratio test comparing the random effects model to a simple OLS model (used to compute linear regression consumption predictions). In each case, the test has a p-value of < 0.001 , rejecting the null hypothesis of no area-level random effects.

4.1 Estimates of local consumption spending

Figures 4.1, 4.2 and 4.3 map our EB estimates of weekly equivalised consumption across local authorities in Great Britain, for each of the three consumption measures. We find that average equivalised weekly deflated consumption estimates (C3) range between £279 per week (in Barking and Dagenham) and £478 (in Chiltern). The median average level of equivalised weekly deflated consumption is £372. Areas in the South East of England, such as Waverley (£468), have relatively high levels of deflated consumption, in contrast to areas such as Sandwell (£294) and the City of Kingston upon Hull (£306) in the West Midlands. Within London, we observe significant variation across local authorities. Barking and Dagenham (£279)

exhibits the lowest level of deflated consumption nationally, whilst residents of the City of London (£451) and Richmond upon Thames (£448) have the fourth and eighth highest level of deflated consumption in the country.

Non-housing consumption (C1) (£)

170 220 270 320 370




Figure 4.1: *EB estimates of 2018/2019 Equivalised Consumption excluding housing-related expenditures (C1)*

We compare the distributions of each consumption measure across and within broad regions more explicitly in Figure 4.4. Several interesting patterns emerge, in particular relating to the ranking of London. On average, London local authorities rank at the top of the consumption distribution when comparing across regions (24.8% above the Great Britain mean) if housing-related expenditures are included. Excluding housing expenditures cuts London's premium to just 0.4% above the national mean. Deflating reduces it further to 2.97% below, implying high nominal housing expenditure in these areas largely reflects high prices. Whether these results imply that households in London enjoy higher or lower living standards than those in other parts of the country depends on the importance of mobility frictions that prevent house prices from perfectly reflecting local amenity flows.¹²

¹²Figures B.1, B.2 and B.3 in the Appendix compare local authority rankings across measures using scatterplots, highlighting local authorities where the ranking changes most.

Consumption including housing (C2) (£)

250 320 390 460 530 600




Figure 4.2: *EB estimates of 2018/2019 Equivalised Consumption including housing-related expenditures (C2)*

Notes: We winsorise the distribution of consumption at £600 to ensure the color scale on the figures are not dominated by outliers.

Deflated consumption including housing (C3) (£)

260 340 420 500




Figure 4.3: *EB estimates of 2018/2019 Equivalised Deflated Consumption (C3)*




Figure 4.4: *Within-region variation (across LAs) in consumption measures*

Notes: Boxes show the interquartile range of equivalised consumption across local authorities within each region. The cross markers show population-weighted regional means. Index is normalised so that the GB mean equals 100 for each measure.

This pattern is mirrored to a lesser extent for local areas in the South East, whereas local authorities in regions such as the North East and Yorkshire and the Humber appear richer when excluding or deflating housing-related expenditures. Inequality – both across and within regions – appears less pronounced when excluding or deflating housing-related expenditures, compared with when nominal expenditures on housing are included.¹³ These results imply that geographic inequality in living standards is more pronounced under the assumptions of the Rosen-Roback model, where high nominal housing expenditure reflects high amenity flows, than under the assumption of substantial mobility frictions, where deflating consumption is more appropriate.

4.2 Standard errors

As outlined in Section 2.3, we obtain area-specific standard errors for our EB estimates, generated via a parametric bootstrap procedure. In Figure 4.5, we plot our estimates of equivalised consumption with the corresponding 95% confidence intervals, ordered by rank in the equivalised consumption distribution. We show results for our measure of total consumption including housing costs, with values deflated according to local house prices (results for alternative consumption measures are qualitatively similar). The size of these confidence intervals varies across areas, reflecting differences in LCFS sample sizes across local authorities. In a few areas, the estimates lie outside the bootstrapped confidence intervals due to biases in our estimation procedure (discussed below). Despite the advantages of the small area estimation procedure, these confidence intervals can be wide in some areas. This means that users should be cautious for example in using these methods to obtain an exact ranking of local authorities. As we show in Section 4.3, however, the EB estimator substantially improves upon the sample mean estimator in terms of variance.

¹³We calculate summary measures of each inequality measure and discuss them in Section 4.4.




Figure 4.5: *EB estimates of 2018/2019 Deflated Equivalised Consumption (C3) with 95% confidence bands*

Notes: We simulate 95% confidence intervals using the bootstrap method in Section 2.3.

4.3 Comparison of estimation methods

In this section, we use the model-based procedure set out in Section 2.2 to obtain results on the statistical performance of the EB estimator relative to the LCFS sample mean, as well as whether inclusion of bank data improves the estimator’s accuracy.¹⁴ Unless otherwise stated, the results in this section pertain to the measure of consumption that includes housing expenditures, and which has been deflated for local housing costs¹⁵.

Figure 4.6 plots the unweighted average (across local authorities) of these mean squared errors across areas split by estimation method. The average MSE of the EB estimator is substantially lower than that of the direct sample mean estimator.¹⁶ Including bank data as covariates in the EB estimation method reduces the mean squared error by 9.2%. This is consistent with our finding (reported in Table A.3

¹⁴We conduct the same analysis using the design-based mean squared error, and report the corresponding results in the Appendix.

¹⁵As shown in Table 4.1, the improvements from including bank data are generally similar for the measure of consumption excluding housing (C1). The relative gains are generally smaller for consumption including housing, likely because bank account data is a less informative predictor of housing-related expenditures.

¹⁶All local authorities with fewer than 10 households in the LCFS are omitted for data disclosure reasons. Given that the performance of the sample mean is likely to be even worse for small areas, it is likely that the difference between the EB estimator and the sample mean would be even greater if including all areas.

in the Appendix) that average spending outgoings from banks accounts is a highly significant predictor of consumption.

In Table 4.1, we report average mean squared errors for each estimation method and consumption method, dividing areas into subgroups based on their LCFS sample size. As expected, the average mean squared error is decreasing in the sample size for all estimators and consumption measures. Both the relative and absolute improvement of the EB estimator over the direct sample mean estimator is greatest for small areas. Those areas with a sample size below 20 also saw the largest absolute reductions in MSE from including bank data in the EB approach, although the percentage improvement is similar across small and large areas.

Sample size bin	Excluding Housing (C1)				Including Housing (C2)				Deflated (C3)			
	EB	EB (no bank)	LR	SM	EB	EB (no bank)	LR	SM	EB	EB (no bank)	LR	SM
30+	187	208	292	953	236	249	689	719	150	166	256	762
20-29	272	297	369	1807	394	415	833	1406	211	233	305	1297
<20	407	443	464	5771	697	721	965	4522	308	339	360	3803

Table 4.1: *Mean MSE by sample size bin, consumption measure, and estimator*

Notes: EB, EB (no bank), LR and SM denote the average mean squared error for the empirical best, empirical best (excluding FINDS data), linear regression and sample mean estimators respectively. Areas with a sample size below 20 are in the bottom half of the sample-size distribution (i.e., the median is 20). Areas with sample sizes of 20-29 households lie between the median and 75th percentile. Areas with sample sizes of 30 or more individuals are in the top quartile.




Figure 4.6: *Average MSE by Estimation Method*

Notes: Results are for average consumption including housing, deflated according to local prices (C3).

In Figure 4.7, we plot area-specific mean squared errors corresponding to the sample mean, linear regression and EB estimators (including the bank account

data), against the LCFS sample size. Consistent with Table 4.1, for all methods we observe that areas with a larger sample size tend to exhibit lower mean squared errors. However, for all areas, the MSE corresponding to the sample mean is higher than the MSE corresponding to the EB estimator. This indicates that the differences in average mean squared error depicted in Figure 4.6 were not just due to very poor performance of the sample mean for the areas with the smallest sample sizes, but rather the EB estimator improves upon the sample mean for all areas (though its relative advantage shrinks for areas with greater sample size).




Figure 4.7: *Area-Level MSE, by Method*

Notes: Results are for average consumption including housing, deflated according to local prices (C3). To avoid potential disclosure, we omit areas with sample sizes less than 10.

To gain intuition for why the EB estimator improves upon the sample mean, we can decompose the expression for the area-specific model MSE in Equation 2.4 as:

$$MSE_{\mathcal{M}}(\hat{c}_a^{EB}) = V_{\mathcal{M}}(\hat{c}_a^{EB} - \bar{c}_a) + \{\mathbb{E}_{\mathcal{M}}(\hat{c}_a^{EB} - \bar{c}_a)\}^2 \quad (4.1)$$

and similarly for the direct estimator (the sample mean):

$$MSE_{\mathcal{M}}(\hat{c}_a) = V_{\mathcal{M}}(\hat{c}_a^{Direct} - \bar{c}_a) + \{\mathbb{E}_{\mathcal{M}}(\hat{c}_a^{Direct} - \bar{c}_a)\}^2 \quad (4.2)$$

Note that because the target parameter $\bar{c}_a$ is a random variable (varying across draws from the superpopulation model 2.3), the usual decomposition of the MSE as the sum of squared bias and variance does not hold. Instead, for each local authority, the mean squared error is the sum of the squared model bias and the variance of the *estimation errors*. We plot this decomposition in Figure 4.8. We observe that intuitively, the variance of the sample mean estimator is very high

for undersampled areas, but as the number of households increases, it appears to converge to that of the EB estimator. The linear regression imputed consumption estimates have an intermediate variance for small areas, but this does not shrink as the local sample size increases (since the estimator does not use any local authority-specific information), implying that for large areas, the variance of this estimator is relatively high.

With regard to the bias, we find that all three estimators are unbiased on average across local authorities. However, for areas with smaller samples, the bias of the sample mean is often non-zero for individual areas. One possibility is that the LCFS is not a representative sample for the APS population (even after using the LCFS weights, as these are stratified by region but not by local authority, [Office for National Statistics \(2023b\)](#)), implying the LCFS sample mean will be a biased estimator of the APS mean.¹⁷ The linear regression estimator exhibits larger absolute biases than the EB estimator for 222 out of 291 reported areas, and the sample mean is more biased than the EB estimator for all but 5 local authorities. As the sample size grows, the bias of the sample mean approaches zero, as expected.




Figure 4.8: *Area-Level Decomposition of MSE, by Method*

Notes: Results are for average consumption including housing, deflated according to local prices (C3). To avoid potential disclosure, we omit areas with sample sizes less than 10.

¹⁷To attempt to correct for this bias, we constructed our own weights by stacking the LCFS and APS datasets, estimating a probit model of belonging to the APS, and then weighing households in the LCFS as $w_i = \frac{1}{\mathbb{P}(i \in \text{LCFS} | X)}$ (to give more weight to the sorts of households that are undersampled in the LCFS compared to the APS). Because we have small or zero sample sizes for some local authorities in the LCFS, we are unable to estimate this probit separately by local area, implying these weights do not perfectly map the LCFS to the APS for each individual area. Perhaps because of this, we found that the use of these weights did not substantially improve the bias in the sample mean for small areas in Figure 4.8.

In the Appendix, we report estimates of the mean-squared errors for our deflated consumption measure using the ‘hold-out sample’ estimator described in Section 2.2. Because many areas are not large enough to permit us to split the sample in two, we restrict attention to areas with an LCFS sample size greater than 20. Despite this being the subsample of areas where we would expect the sample mean to perform best, Table B.7 shows the EB estimator continues to outperform the sample mean and linear regression estimators. Using the hold-out sample estimator, we find that the inclusion of bank account data yields a smaller improvement in MSE (2.3%) compared with the model-based estimator. This is likely due to the omission of areas with fewer than 20 households (necessary to ensure a hold-out sample exists), where we might expect the inclusion of local spending data to have a larger benefit in estimating consumption.

4.4 Comparing local consumption and income

To understand how the geographical distribution of consumption differs from that of income, in this section we compare our estimates of area-level consumption to local authority income averages derived from National Accounts Gross Disposable Household Income (GDHI) data. We adjust the GDHI measure to better reflect cash income as usually understood (following Judge and McCurdy (2022)), as described in Section 3.¹⁸

There are a number of reasons why we would expect these measures of average incomes to differ from our consumption measures. One set of reasons relate to issues of measurement. Firstly, the income data is per capita while the consumption measure is an equivalised value for the household. Since the population data used for GDHI per capita measures includes children and retirees, we would expect these income figures to be lower than average household equivalised area level income. Secondly, there are also reasons to think (unequivalised) consumption might be understated relative to income in this comparison, as grossed up totals from the LCFS and other consumption surveys tend to be lower than implied by other sources (such as the national accounts, Barrett et al. (2015)). Thirdly, the income measure, being based on administrative data, may be more likely to capture very high income households that might be less likely to respond to the consumption survey than

¹⁸Though the GDHI is our preferred measure of income, we also estimate mean incomes using the LCFS (and the EB estimator). We find our income estimates are highly correlated with the GDHI measure of income, and report correlations between these (and the measures of consumption) in Table B.2. EB-estimated income is better correlated with each of our consumption measures than the GDHI, suggesting some of the differences between GDHI-measured income and consumption result from differences due to the LCFS sample. However, the improvement is marginal, suggesting the patterns are not entirely driven by sampling differences.

others. Since incomes tend to be highly skewed, the presence of even only a few high income individuals could potentially have a strong influence on local average incomes, while not having as great an effect on our consumption measure.¹⁹

A second class of reasons for differences across income and consumption measures relate to economic incentives and behaviour. People tend to live in different sorts of locations over their lives, and these may coincide with ages when incomes tend to be higher or lower than average. This would drive a wedge between income and consumption that is greater in some local authorities than others. As noted in the introduction, this means there are strong theoretical reasons to believe that household consumption is a better measure of lifetime resources and living standards than current period incomes. For example, people might live in London at ages when incomes and savings tend to peak, but then retire to other locations to draw down their savings. Another possible reason is that savings behaviour differs across locations even at a given stage of life. For example relative price differences across areas could make it advantageous for those planning to leave expensive urban areas to postpone their spending. Average preferences could also in principle differ across areas.

We plot the distribution of weekly equivalised deflated consumption against the mean per-capita income distribution, along with a linear fit, in Figure 4.9. Of particular note is the disparity between the position of London areas in the income and consumption distributions. For example, Tower Hamlets ranks at the 93rd percentile of the local authority distribution of income per capita, but at the 2nd percentile of the equivalised deflated consumption distribution. Of the ten areas with the largest absolute difference between their ranks on measures of income and deflated consumption, all are located in London. We plot the comparison with consumption excluding housing in Figure B.4, and report pairwise rank correlations between all our measures of consumption and income in Table B.2.

¹⁹In Figure B.5 in the Appendix, we compare our relative consumption measures to the local authority ranking according to the 2019 Income Index of Multiple Deprivation (IMD) for English local authorities. Since this measure is based on the share of households below a given threshold of income (see Section 3), it is less likely to be affected by outliers in the right tail of the income distribution.




Figure 4.9: *Equivalised Deflated Consumption including Housing (C3) vs per-capita income - ranks*

Notes: The local authority classification in our income and consumption data differs (as the boundaries have changed over time), implying 28 local authorities do not merge immediately between the datasets. For these, we construct a mapping manually, displayed in Table B.3.

We also calculate percentile ratios from the distribution of equivalised consumption. We find the local authority at the 90th percentile of the income distribution earns (on average) 65% more than the local authority at the 10th percentile, whereas the corresponding local authority gap for deflated consumption is 27%. This qualitative pattern holds for all the ratios reported in Table 4.2, and is starkest at the extremes: individuals in the richest local authority (City of London) have mean incomes 16.2 times higher than those in the poorest local authority (Leicester), but only consume 1.72 times as much. These results are consistent with the theoretical (Friedman, 1957) and empirical (Meyer and Sullivan, 2023) literature arguing that consumption is less unequally distributed than income. However, the consumption measure including housing-related expenditure (and imputing consumption flows from housing for owners) exhibits similar levels of inequality (measured using percentile ratios) as income (though excluding housing expenditures reduces the degree of inequality). Additionally, the same remarks regarding measurement apply: most notably our consumption measures are adjusted with an equivalence scale whereas the income measures are per capita.

		Richest vs Poorest	95:5	90:10	75:25
1	Consumption excluding housing (C1)	2.07	1.56	1.44	1.21
2	Consumption including housing (C2)	2.78	1.87	1.67	1.33
3	Deflated Mean Consumption (C3)	1.72	1.35	1.27	1.13
4	Mean Income per Capita	16.18	1.84	1.65	1.27

Table 4.2: *Percentile Ratios of local authorities*

5 Conclusion

This paper presents estimates of average equivalised consumption across local authority districts for Great Britain combining data from a household budget survey, a much larger population survey and area-level predictors of consumption spending.

Bank account data, and other measures of ‘financial footprints’ data are increasing in their availability. By themselves, such measures can provide useful indicators of living standards in different areas. But they fall short of giving a complete picture of households’ consumption levels. Our findings suggest that financial footprints data can however help to substantially improve estimates of consumption derived from household budget surveys. There are thus good reasons to hope that accurate area-level estimates of consumption spending will soon be available to complement existing income-based measures of local living standards.

References

- Barrett, G., P. Levell, and K. Milligan (2015). A Comparison of Micro and Macro Expenditure Measures across Countries Using Differing Survey Methods. In C. D. Carroll, T. F. Crossley, and J. Sabelhaus (Eds.), *Improving the Measurement of Consumer Expenditures*, pp. 263–286. University of Chicago Press.
- Battese, G. E., R. M. Harter, and W. A. Fuller (1988). An Error-Components Model for Prediction of County Crop Areas Using Survey and Satellite Data. *Journal of the American Statistical Association* 83(401), 28–36.
- Benson, R., B. Duffy, R. Hesketh, and K. Hewlett (2024). Attitudes to Inequalities. *Oxford Open Economics* 3(Supplement_1), i39–i63.
- Blundell, R. and I. Preston (1995). Income, Expenditure and the Living Standards of UK Households. *Fiscal Studies* 16(3), 40–54.

- Brewer, M., B. Etheridge, and C. O’Dea (2017, October). Why are Households that Report the Lowest Incomes So Well-Off? *The Economic Journal* 127(605), F24–F49.
- Bryan, G. and M. Morten (2019). The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia. *Journal of Political Economy* 127(5), 2229–2268.
- Deaton, A. and S. Zaidi (2002, May). Guidelines for Constructing Consumption Aggregates for Welfare Analysis. World Bank Publications - Reports 14101, The World Bank Group.
- Diamond, R. and E. Moretti (2021, December). Where is Standard of Living the Highest? Local Prices and the Geography of Consumption. Working Paper 29533, National Bureau of Economic Research. Revision date: August 2024.
- Fay, R. E. and R. A. Herriot (1979). Estimates of income for Small Places: An application of James-Stein Procedures to Census Data. *Journal of the American Statistical Association* 74(366), 269–277.
- Friedman, M. (1957). *Theory of the Consumption Function*. Princeton University Press.
- González-Manteiga, W., M. J. Lombardía, I. Molina, D. Morales, and L. Santamaría (2008). Analytic and Bootstrap Approximations of Prediction Errors under a Multivariate Fay–Herriot Model. *Computational Statistics & Data Analysis* 52(12), 5242–5252.
- Judge, L. and C. McCurdy (2022). Income Outcomes. Technical report, Resolution Foundation.
- Kennan, J. and J. R. Walker (2011). The Effect of Expected Income on Individual Migration Decisions. *Econometrica* 79(1), 211–251.
- McCann, P. (2020). Perceptions of Regional Inequality and the Geography of Discontent: Insights from the UK. *Regional Studies* 54(2), 256–267.
- McLennan, D., S. Noble, M. Noble, E. Plunkett, G. Wright, and N. Gutacker (2019). English Indices of Deprivation 2019: Technical Report. Technical report, Ministry of Housing, Communities and Local Government.
- Meyer, B. D. and J. X. Sullivan (2023). Consumption and Income Inequality in the United States since the 1960s. *Journal of Political Economy* 131(2), 247–284.
- Molina, I. and J. N. K. Rao (2010). Small Area Estimation of Poverty Indicators. *The Canadian Journal of Statistics / La Revue Canadienne de Statistique* 38(3), 369–385.
- Office for National Statistics (2018). Regional Gross Disposable Household Income QMI. Technical report, Office for National Statistics.

- Office for National Statistics (2019). Price Index of Private Rents, UK: monthly price statistics. Dataset.
- Office for National Statistics (2021). Living Costs and Food Survey - UK. ONS SRS Metadata Catalogue. Dataset, released 27 October 2021.
- Office for National Statistics (2023a, January). Housing, England and Wales: Census 2021. Statistical bulletin. Release date: 5 January 2023.
- Office for National Statistics (2023b, August). Living Costs and Food Survey QMI. Technical report, Office for National Statistics.
- Office for National Statistics (2024). Regional Gross Disposable Household Income, UK: 1997 to 2022. Technical report.
- Office for National Statistics (2025). Annual Population Survey - UK. ONS SRS Metadata Catalogue. Dataset, released 24 October 2025.
- Pfeffermann, D. (2013). New Important Developments in Small Area Estimation. *Statistical Science* 28(1), 40–68.
- Poterba, J. M. (1989). Lifetime Incidence and the Distributional Burden of Excise Taxes. *The American Economic Review* 79(2), 325–330.
- Roback, J. (1982). Wages, Rents, and the Quality of Life. *Journal of political Economy* 90(6), 1257–1278.
- Rosen, S. (1979). Wage-based indexes of urban quality of life. *Current issues in urban economics*, 74–104.
- Samuelson, P. A. (1974, December). Complementarity-An Essay on the 40th Anniversary of the Hicks-Allen Revolution in Demand Theory. *Journal of Economic Literature* 12(4), 1255–1289.
- Slesnick, D. T. (1993). Gaining Ground: Poverty in the Postwar United States. *Journal of Political Economy* 101(1), 1–38.

Appendix A Regression results

Table A.1: *Regression Results: consumption excluding housing (C1)*

Variable	Estimate	Std. Error	t value
Intercept	-0.367	0.322	-1.142
Dummy: Household contains an U14 individual	0.050	0.033	1.512
Dummy: Household contains an individual aged 14-18	0.047	0.068	0.686
Dummy: Household contains >1	0.444	0.023	18.996
Number of U14s in household	0.027	0.017	1.602
Number of 14-18 in household	0.104	0.055	1.905
Number of additional adults	0.229	0.014	15.899
Sex of household head	0.008	0.014	0.599
MSOA-level log of average card spending (from Natwest data)	0.337	0.031	10.911
Dummy: household head finished education between 18 and 21	0.145	0.021	7.058
Dummy: household head finished education aged over 21	0.188	0.021	9.154
Dummy: household head is still in education	0.273	0.065	4.183
Ethnicity: White	0.180	0.045	4.034
Ethnicity: Asian/Asian British	-0.143	0.055	-2.612
Ethnicity: Black African/Caribbean/Black British	-0.177	0.065	-2.735
Dummy: Household Head is aged 30-40	0.020	0.030	0.669
Dummy: Household Head is aged 40-50	0.042	0.030	1.377
Dummy: Household Head is aged 50-60	-0.010	0.031	-0.333
Dummy: Household Head is aged 60-70	-0.086	0.033	-2.633
Dummy: Household Head is aged 70+	-0.259	0.034	-7.710
HH is mortgager	0.526	0.018	29.793
HH is owner	0.585	0.019	30.131
LAD-median electricity consumption (in log KWH)	0.018	0.021	0.861
Log median house price in local authority	0.124	0.020	6.193
Dummy: observation is from 2018	-0.027	0.014	-2.018
Individual error variance	0.362		
Random error variance	0.004		
LR test of random effects model (p-value)	11.064 (<0.001)		
Sample size	8744		

Notes: These parameter estimates correspond to a random effects regression of log consumption excluding housing-related expenditure on covariates and a local-authority random effect. It was estimated on a sample of 8744 households from the 2018 and 2019 Living Costs and Food Survey. The LR test is the likelihood ratio test statistic obtained from testing the random effects model against a linear model omitting the random area effect. The omitted reference ethnicity is ‘Other ethnic group’, and the omitted age of household head group is ‘Household Head is aged 18-30’.

Table A.2: *Regression Results: consumption including housing (C2)*

Variable	Estimate	Std. Error	t value
Intercept	-1.065	0.214	-4.982
Dummy: Household contains an U14 individual	0.034	0.019	1.751
Dummy: Household contains an individual aged 14-18	0.006	0.040	0.139
Dummy: Household contains >1 adult	0.235	0.014	17.278
Number of U14s in household	0.026	0.010	2.640
Number of 14-18 in household	0.089	0.032	2.786
Number of additional adults	0.163	0.008	19.463
Sex of household head	0.002	0.008	0.203
MSOA-level log of average card spending (from Natwest data)	0.199	0.019	10.720
Dummy: household head finished education between 18 and 21	0.094	0.012	7.853
Dummy: household head finished education aged over 21	0.123	0.012	10.283
Dummy: household head is still in education	0.177	0.038	4.644
Ethnicity: White	0.089	0.026	3.410
Ethnicity: Asian/Asian British	-0.099	0.032	-3.115
Ethnicity: Black African/Caribbean/Black British	-0.076	0.038	-2.020
Dummy: Household Head is aged 30-40	0.034	0.018	1.930
Dummy: Household Head is aged 40-50	0.064	0.018	3.628
Dummy: Household Head is aged 50-60	0.057	0.018	3.159
Dummy: Household Head is aged 60-70	0.026	0.019	1.388
Dummy: Household Head is aged 70+	-0.082	0.020	-4.197
HH is mortgager	0.363	0.010	35.228
HH is owner	0.393	0.011	34.688
LAD-median electricity consumption (in log KWH)	-0.003	0.013	-0.196
Log median house price in local authority	0.372	0.014	26.478
Dummy: observation is from 2018	-0.011	0.008	-1.329
Individual error variance	0.121		
Random error variance	0.005		
LR test of random effects model (p-value)	98.94 (<0.001)		
Sample size	8744		

Notes: These parameter estimates correspond to a random effects regression of log consumption including housing-related expenditure on covariates and a local-authority random effect. It was estimated on a sample of 8744 households from the 2018 and 2019 Living Costs and Food Survey. The LR test is the likelihood ratio test statistic obtained from testing the random effects model against a linear model omitting the random area effect. The omitted reference ethnicity is ‘Other ethnic group’, and the omitted age of household head group is ‘Household Head is aged 18-30’.

Table A.3: *Regression Results: deflated consumption including housing (C3)*

Variable	Estimate	Std. Error	t value
Intercept	2.362	0.189	12.463
Dummy: Household contains an U14 individual	0.034	0.019	1.78
Dummy: Household contains an individual aged 14-18	0.005	0.04	0.123
Dummy: Household contains >1 adult	0.239	0.014	17.647
Number of U14s in household	0.027	0.01	2.708
Number of 14-18 in household	0.092	0.032	2.893
Number of additional adults	0.161	0.008	19.252
Sex of household head	0.003	0.008	0.36
MSOA-level log of average card spending (from Natwest data)	0.217	0.018	12.057
Dummy: household head finished education between 18 and 21	0.092	0.012	7.725
Dummy: household head finished education aged over 21	0.12	0.012	10.04
Dummy: household head is still in education	0.177	0.038	4.67
Ethnicity: White	0.109	0.026	4.187
Ethnicity: Asian/Asian British	-0.100	0.032	-3.152
Ethnicity: Black African/Caribbean/Black British	-0.076	0.037	-2.021
Dummy: Household Head is aged 30-40	0.035	0.017	2.001
Dummy: Household Head is aged 40-50	0.064	0.018	3.623
Dummy: Household Head is aged 50-60	0.058	0.018	3.257
Dummy: Household Head is aged 60-70	0.027	0.019	1.444
Dummy: Household Head is aged 70+	-0.082	0.019	-4.194
HH is mortgager	0.365	0.01	35.642
HH is owner	0.397	0.011	35.205
LAD-median electricity consumption (in log KWH)	0.026	0.012	2.131
Log median house price in local authority	0.071	0.012	5.996
Dummy: observation is from 2018	-0.021	0.008	-2.608
Individual error variance	0.121		
Random error variance	0.002		
LR test of random effects model (p-value)	16.569 (<0.001)		
Sample size	8744		

Notes: These parameter estimates correspond to a random effects regression of log deflated consumption on covariates and a local-authority random effect. It was estimated on a sample of 8744 households from the 2018 and 2019 Living Costs and Food Survey. The LR test is the likelihood ratio test statistic obtained from testing the random effects model against a linear model omitting the random area effect. The omitted reference ethnicity is ‘Other ethnic group’, and the omitted age of household head group is ‘Household Head is aged 18-30’.

Appendix B Additional Tables and Figures

Table B.1: *Covariate Balance*

Variable	LCFS Sample Mean	APS Sample Mean
Number of U14s in household	0.42	0.41
Number of 14-18 in household	0.10	0.11
Number of additional adults	0.84	0.85
Sex of household head	0.61	0.59
Share in social housing	0.15	0.17
Ethnicity: White	0.92	0.91
Ethnicity: Asian/Asian British	0.04	0.04
Ethnicity: Black African/Caribbean/Black British	0.02	0.02
Ethnicity: Other Ethnic Group	0.02	0.02
Dummy: HH Head is aged 18-30	0.07	0.07
Dummy: HH Head is aged 30-40	0.16	0.15
Dummy: HH Head is aged 40-50	0.17	0.17
Dummy: HH Head is aged 50-60	0.19	0.2
Dummy: HH Head is aged 60-70	0.18	0.18
Dummy: HH Head is aged 70+	0.22	0.23
Dummy: HH is mortgager	0.32	0.3
Dummy: HH is owner	0.38	0.38
Dummy: HH is renter	0.3	0.32
LAD-median electricity consumption (in log KWH)	2.96	2.9
Share of sample who are renters in non social housing	0.15	0.15
Log median house price in local authority	12.28	12.22
-level log of average card spending (from Natwest data)	10.18	10.15
Dummy: HH Head finished education under 18	0.57	0.59
Dummy: HH Head finished education between 18 and 21	0.2	0.2
Dummy: HH Head finished education aged over 21	0.21	0.2
Dummy: HH Head still in education	0.02	0.01

Notes: This table contains mean values of covariates in the APS and LCFS samples used when estimating area-level mean consumption.




Figure B.1: *Consumption excluding housing (C1) versus consumption including housing-related expenditures (C2)*

Notes: We label the five local authorities with the largest absolute disparity in rank between the two measures.




Figure B.2: *Consumption excluding housing (C1) versus deflated consumption (C3)*

Notes: We label the five local authorities with the largest absolute disparity in rank between the two measures.




Figure B.3: *Consumption including housing (C2) versus deflated consumption (C3)*

Notes: We label the five local authorities with the largest absolute disparity in rank between the two measures.

	C1	C2	C3	Income IMD	GDHI Income	EB-estimated Income
C1	1.000	0.845	0.975	0.882	0.802	0.880
C2	0.845	1.000	0.776	0.687	0.912	0.935
C3	0.975	0.776	1.000	0.889	0.741	0.826
Income IMD	0.882	0.687	0.889	1.000	0.701	0.770
GDHI Income	0.802	0.912	0.741	0.701	1.000	0.903
EB-estimated Income	0.880	0.935	0.826	0.770	0.903	1.000

Table B.2: *Spearman Correlation Matrix of local authorities by consumption, Income IMD, EB-estimated income and GDHI Income.*

Notes: C1 is consumption excluding housing expenditures, C2 includes nominal housing expenditures, C3 is deflated using estimates of local housing costs. The Index of Multiple Deprivation is only available for English local authorities, so all correlations are based on the subsample of English local authorities only.




Figure B.4: *Equivalised Consumption excluding housing (C1) vs per-capita income - ranks*

Notes: We label the five local authorities with the largest absolute disparity in rank between the two measures.




Figure B.5: *Equivalised Consumption excluding housing (C1) vs Income IMD - ranks*

Notes: The Index of Multiple Deprivation is only available for English local authorities. We label the four local authorities with the largest absolute disparity in rank between the two measures.




Figure B.6: *GDHI Income vs Income IMD - ranks*

Notes: The Index of Multiple Deprivation is only available for English local authorities. We label the four local authorities with the largest absolute disparity in rank between the two measures.




Figure B.7: *Average area-level MSE by method: hold-out sample estimator*

Notes: Results are for average consumption including housing, deflated according to local prices (C3).




Figure B.8: *Area-Level MSE, by Method - hold-out sample estimator*

Notes: Results are for average consumption including housing, deflated according to local prices (C3). To avoid potential disclosure, we omit areas with sample sizes less than 10.




Figure B.9: *Area-Level Decomposition of MSE, by Method - hold-out sample estimator*

Notes: Results are for average consumption including housing, deflated according to local prices (C3). To avoid potential disclosure, we omit areas with sample sizes less than 10.

Table B.3: *Local authority mapping for unmerged areas*

	LA in Consumption Data	LA in Income data
1	Richmondshire	North Yorkshire
2	Craven	North Yorkshire
3	Corby	North Northamptonshire
4	Ryedale	North Yorkshire
5	Wellingborough	North Northamptonshire
6	Daventry	West Northamptonshire
7	Kettering	North Northamptonshire
8	South Bucks	Buckinghamshire
9	East Northamptonshire	North Northamptonshire
10	Eden	Westmorland And Furness
11	South Northamptonshire	West Northamptonshire
12	Selby	North Yorkshire
13	Hambleton	North Yorkshire
14	Barrow-In-Furness	Westmorland And Furness
15	Copeland	Cumberland
16	Scarborough	North Yorkshire
17	Chiltern	Buckinghamshire
18	Mendip	Somerset
19	Allerdale	Cumberland
20	Sedgemoor	Somerset
21	South Lakeland	Westmorland And Furness
22	Carlisle	Cumberland
23	Harrogate	North Yorkshire
24	Northampton	West Northamptonshire
25	South Somerset	Somerset
26	Somerset West And Taunton	Somerset
27	Aylesbury Vale	Buckinghamshire
28	Wycombe	Buckinghamshire