Estimating local consumption spending by combining household surveys and financial footprints data

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

August, 2025

Abstract

We estimate average equivalised consumption measures across 367 local authority districts in Great Britain. 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 card transactions and energy consumption. Simulations indicate that our measures significantly outperform sample averages taken from the budget survey alone, as well as naive regression imputation estimates. We also find that including transactions data substantially improves our estimates, suggesting these data could play an important role in measuring local consumption and hence living standards in the future. We compare our consumption estimates to local income measures and show the former is less unequally distributed across areas, and that the ranking of local authorities in terms of living standards is different under the two measures.

Keywords: Small area estimation, imputation, local consumption

JEL classification: E21, I32, R12

Acknowledgements: This output is for 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/M010147/1) is also gratefully acknowledged. All errors and omissions remain the responsibility of the authors. This project makes use of data on credit and debit card transactions kindly shared with the IFS from the company Fable. Fable bears no responsibility for any errors in the interpretation of their data and the views expressed in this paper are our own. Correspondence: peter_l@ifs.org.uk, gautam.vyas@ifs.org.uk

[*]Institute for Fiscal Studies and ESCoE

[†]Institute for Fiscal Studies, UCL, 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 in different countries using 28 different measures, based on regional disposable income, or output or 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.

In this paper, we produce estimates of resident households' average *consumption* spending for 367 local authority districts across Great Britain. Because there is no survey of household consumption spending with sufficient sample size to precisely measure average in all local authorities, we use small area estimation methods to combine 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 local households' credit and debit card spending, and household energy use. We evaluate these methods using simulations and show that they 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.

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). One recent article by Meyer and Sullivan (2023) puts the 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]

One reason that incomes are a more popular measure of local living standards is that - for small areas in particular, large-scale sources of data are needed to calculate precise averages at a local level. While information on individual incomes is increasingly available from tax records there are no similar large-sample measures of local consumption. Rather, consumption is still largely measured in relatively small surveys (for example, the Living Costs and Food Survey, LCFS, includes just over 5,000 households each year), which may have small or zero samples in some spatial units. This means that existing estimates tend to only produce estimates for relatively aggregated regions (Office for National Statistics, 2020). While private sector sources of bank data are becoming increasingly common, these are not yet available to many national statistical authorities, and can be selective in their coverage. Moreover, they do not typically include demographic information needed to equivalise consumption values according to household needs.

The literature on *small area estimation* - or methods to estimate statistical objects of interest (such as means, counts or proportions) in areas where sample sizes are small or non-existent (for a survey see Pfeffermann (2013)) - suggests approaches that can be used to most effectively combine information that is available from the different sources at our disposal. We use the 'empirical best' estimator set out in Molina and Rao (2010), which in effect combines parametric and non-parametric/'direct' estimates of local consumption - with the latter given greater weight in areas with a larger sample size. We show that this method greatly outperforms direct estimates of local means in terms of mean-squared error. We also show that the addition of area-level covariates from card transaction data improves the quality of the estimates. The growing availability of 'financial footprints' data could therefore be used to improve the accuracy of local consumption measures if national statistical agencies chose to adopt these methods.

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 consumption.

[1] Another 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, **?**).

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 sets our results, including mean average consumption according to different measures and standard errors around our estimates. Section 5 concludes.

2 Small area estimation of means

Suppose we wish to estimate the mean of some variable $y_{i,a}$ among individuals i in an area a (across A different areas).[2] Let the population in each area be given by N_a such that $\bar{y}_a = \sum \frac{y_{i,a}}{N_a}$. However, suppose we only observe a sample s (with size $n_a \leq N_a$) with observations on $y_{i,a}$ in a survey dataset.

The most straightforward estimator for $\bar{y}_a$ is the so-called the *direct* estimator calculated using the survey data alone

$$\hat{\bar{y}}_a^{Direct} = \sum_i^{n_a} \frac{y_{i,a}}{n_a} \tag{2.1}$$

While this estimator is unbiased for the population mean (assuming random sampling), 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. This additional data could be a vector of covariates $\mathbf{X}_i$ for individuals in each area from a population census or another survey (in which the variable $\mathbf{y}$ is not observed), and/or area-level information on predictors of $\mathbf{y}$ from some other dataset, $\mathbf{Z}_a$ (denoting the rowstacks of these matrices across individual observations as $\mathbf{X}$ and $\mathbf{Z}$).

Suppose we partition the vector of observations on $\mathbf{y}$ into two subvectors $\mathbf{y} = (\mathbf{y}_s', \mathbf{y}_r')'$, where $\mathbf{y}_s$ comprises the observations that are included in sample s and $\mathbf{y}_r$ is made up of the out-of-sample observations we wish to predict, and for which we observe $\mathbf{X}$ from a census or other large survey. If the matrix of covariates $\mathbf{X}$ are available in the survey dataset and so can be similarly partitioned into $\mathbf{X}_s$ and $\mathbf{X}_r$, then we can estimate a model of $\mathbf{y}_s|\mathbf{X}_s, \mathbf{Z}$

$$y_{i,a} = f(\mathbf{X}_i, \mathbf{Z}_a, \varepsilon_i) \tag{2.2}$$

[2] To present the main idea in its simplest form, we focus on the mean. The small area approach can be extended to estimate other statistics like the variance or quantiles - see Molina and Rao (2010) for more details.

and use this to impute the values of $\mathbf{y}_r$. This allows us to calculate an estimated mean value of y in each area as

$$\tilde{\bar{y}}_a = \frac{1}{n_a}\sum_{i\in s} y_{i,a} + \frac{1}{N_a - n_a}\sum_{i\in r} E[f(\mathbf{X}_i, \mathbf{Z}_a, \varepsilon_i)|\mathbf{y}_{a,s}] \tag{2.3}$$

The challenge is to impute the missing values of y using the best possible out-of-sample predictor. *Empirical best predictors* are approaches that do this in a way that minimise the mean squared error (MSE) of the resulting estimates $E\left[(\tilde{\bar{y}}_a - \bar{y}_a)^2\right]$. Intuitively, we want to account for area level differences in the model residuals ε as well as the distributions of local covariates $\mathbf{X}$ and $\mathbf{Z}$, but not rely too heavily on estimates of the average residuals from the survey data that are noisy and unreliable.

A method to address this challenge is perhaps easiest to understand when outcomes are normally distributed. Suppose $y_{i,a} \sim N(\boldsymbol{\mu}_a, \mathbf{V}_a)$ with $\boldsymbol{\mu}_a = (\boldsymbol{\mu}_r', \boldsymbol{\mu}_s')'$ and

$$\mathbf{V}_a = \left(\begin{bmatrix} \mathbf{V}_{a,s} & \mathbf{V}_{a,sr} \\ \mathbf{V}_{a,rs} & \mathbf{V}_{a,r} \end{bmatrix}\right)$$

The best predictors for outcomes for out of sample observations are the conditional means $\boldsymbol{\mu}_{a,r|s}$. By the properties of multivariate normal distributions, the conditional distribution of observations in $\mathbf{y}_r$ in area a given $\mathbf{y}_s$ in area a will be

$$\mathbf{y}_{a,r|s} \sim N(\boldsymbol{\mu}_{a,r|s}, \mathbf{V}_{a,r|s}) \tag{2.4}$$

where

$$\boldsymbol{\mu}_{a,r|s} = \boldsymbol{\mu}_{a,r} + \mathbf{V}_{a,rs}\mathbf{V}_{a,s}^{-1}\left(\mathbf{y}_{a,s} - \boldsymbol{\mu}_{a,s}\right) \tag{2.5}$$

and

$$\mathbf{V}_{a,r|s} = \mathbf{V}_{a,r} - \mathbf{V}_{a,rs}\mathbf{V}_{a,s}^{-1}\mathbf{V}_{a,sr} \tag{2.6}$$

The formula for the conditional mean 2.5 shows how estimates of the average model residuals in area a, $(\mathbf{y}_{a,s} - \boldsymbol{\mu}_{a,s})$, should be scaled down according to the sample variance, and scaled up according to the covariance between values in r and s. This covariance can be non-zero if residuals are not independently and identically distributed. This can arise in random effect models such as the nested error linear regression model (Battese et al., 1988) as we now explain.

2.1 Nested error linear regression model for log consumption

In what follows, we will assume a nested error linear regression model for log consumption of the form

$$\begin{aligned} \log c_{i,a} &= \mathbf{X}_{i,a}\boldsymbol{\beta} + \mathbf{Z}_a\boldsymbol{\pi} + u_a + \varepsilon_{i,a} \\ u_a &\sim_{iid} N(0, \sigma_u^2) \\ \varepsilon_{i,a} &\sim_{iid} N(0, \sigma_\varepsilon^2) \end{aligned} \tag{2.7}$$

This model includes random area-specific residuals as well as individual level errors. This version also includes both individual-level and area-level covariates.[3] Stacking the elements for each area, define the $N_a \times 1$ vector of log consumption and individual errors for individuals in area a as $\log \mathbf{c}_a = [\log c_{1,a}; ...; \log c_{N_a,a}]$ and $\boldsymbol{\varepsilon}_a = [\varepsilon_{1,a}; ...; \varepsilon_{N_a,a}]$. Similarly, define the $N_a \times K$ matrix of covariates for area a as $\mathbf{X}_a = [\mathbf{X}_{i,a}; ...; \mathbf{X}_{N_a,a}]$ (where K is the number of columns of $\mathbf{X_{i,a}}$).

The model itself is for a *superpopulation*, which gives rise to many different populations (specific draws of $\mathbf{u}_a$ and $\boldsymbol{\varepsilon}_a$). The A different values of u_a (forming the vector $\mathbf{u}_a$) here are hyper-parameters that can vary across populations.

Molina and Rao (2010) derive the conditional mean and variances for this model, which are particular cases of the more general expressions in 2.5 and 2.6. These are

$$\mu_{a,r|s} = \mathbf{X}_{a,r}\boldsymbol{\beta} + \mathbf{Z}_a\pi + \sigma_u^2 \mathbf{1}_{N_a-n_a}\mathbf{1}'_{N_a-n_a}\mathbf{V}_{a,s}^{-1}\left(\mathbf{y}_{a,s} - \mathbf{X}_{a,s}\boldsymbol{\beta}\right) \tag{2.8}$$

$$\mathbf{V}_{a,r|s} = \sigma_u^2\left(1-\gamma_a\right)\mathbf{1}_{N_a-n_a}\mathbf{1}'_{N_a-n_a} + \sigma_\varepsilon^2\mathbf{I}_{N_a-n_a} \tag{2.9}$$

where $\mathbf{V}_{a,s} = \sigma_u^2\mathbf{1}_{n_d}\mathbf{1}'_{n_d} + \sigma_\varepsilon^2\mathbf{I}_{n_a}$ and $\gamma_a = \sigma_u^2\left(\sigma_u^2 + \sigma_\epsilon^2/n_a\right)^{-1}$.

These expressions allow us to obtain the conditional mean of *log* consumption for each area. This of course is not the same as the conditional mean of consumption. To obtain values for mean consumption, we follow Molina and Rao (2010) in using a Monte-Carlo approximation. This involves drawing L vectors of $\log \mathbf{c}_{a,r}$ from the normal distribution with area-specific conditional mean and variance, exponentiating, and then averaging. To avoid drawing large samples from a multivariate normal distribution (which is computationally intensive), they instead use draws from the model

[3] It is standard to divide small area estimation methods into those that use individual or unit-level covariates, such as (Battese et al., 1988), and those that use area-level information, such as the area-level random effects model of Fay and Herriot (1979) In our approach we make use of both simultaneously.

$$\log \mathbf{c}_{a,r} = \boldsymbol{\mu}_{a,r|s} + v_a \mathbf{1}_{N_a - n_a} + \boldsymbol{\varepsilon}_{a,r} \tag{2.10}$$

where $v_a \sim N\left(0, \sigma_u^2\left(1-\gamma_d\right)\right)$ and $\boldsymbol{\varepsilon}_{a,r} \sim N\left(\mathbf{0}_{N_a-n_a}, \sigma_\varepsilon^2 \mathbf{I}_{N_a-n_a}\right)$. This has an identical variance covariance but only requires taking draws from two independent univariate distributions.

Draws from this model can be then be transformed into imputed values of $\hat{c}_{i,a}$ by taking $exp(\widehat{\log c_{i,a}})$. Calculating average consumption $\left(\hat{\bar{c}}_a^{EB} = \sum_i^{N_a} \frac{\hat{c}_{i,a}}{\Omega(X_i)}\right)$ in each area - where $\Omega(X_i)$ is an equivalisation scale - is then straightforward using equation 2.3.

In what follows, we refer to this procedure as the Empirical Best (EB) estimator, because it minimises the MSE under the assumed model 2.7. In cases, where there are no observations in the sample data, the EB estimator reduces to a *synthetic* estimator, which is a regression prediction based only on $\mathbf{X}_{a,r}$, $\mathbf{Z}_a$ and draws of $\boldsymbol{\varepsilon}_a$.

2.2 MSE estimation

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

$$MSE(\hat{\bar{c}}_a^{EB}) = E\left(\hat{\bar{c}}_a^{EB} - \bar{c}_a^{EB}\right)^2 \tag{2.11}$$

where the expectation E is taken over realisations of $\mathbf{u}_a$ and $\boldsymbol{\varepsilon}_a$. To estimate the MSE's associated with both the EB method and alternative estimation strategies, we use a parametric bootstrap following the steps outlined in Molina and Rao (2010) (who follow the bootstrap method of González-Manteiga et al. (2008)). These are:

1. Estimate the model in 2.7 using maximum likelihood in our consumption sample. Obtain estimators of $\boldsymbol{\beta}$,$\boldsymbol{\pi}$,σ_u^2,σ_ε^2, denoted: $\hat{\boldsymbol{\beta}}$, $\hat{\boldsymbol{\pi}}$, $\hat{\sigma}_u$, $\hat{\sigma}_\varepsilon$.

2. Take draws of the random effects and individual level residuals $u_a^* \sim N(0, \hat{\sigma}_u^2)$, $\varepsilon_{i,a}^* \sim N(0, \hat{\sigma}_\varepsilon^2)$.

3. Use this to construct the bootstrap super-population model

$$\log c_{i,a}^* = \mathbf{X}_{i,a}\hat{\boldsymbol{\beta}} + \mathbf{Z}_a\hat{\boldsymbol{\pi}} + u_a^* + \varepsilon_{i,d}^* \tag{2.12}$$

4. From this bootstrap super-population model, draw B i.i.d bootstrap populations $\{\log c_{i,a}^{*,(b)}; i = 1, ..., N_a, a = 1, .., A\}$, and for each of these populations calculate the parameters $\bar{c}_a^{*(b)} = \frac{1}{N_a}\sum_{i=1}^{N_a} \exp(\log c_{i,a}^{*(b)})/\Omega(X_i)$.

5. For each bootstrap population, take values from with the same indices as the initial sample and use these to produce the EB estimators as described

in Section 2.1, synthetic regression estimators described in 3.2 and direct estimators, and a version of the EB estimator where we omit data on $\mathbf{Z}_a$. Denote these $\hat{\bar{c}}_a^{EB*(b)}$, $\hat{\bar{c}}_a^{Synthetic*(b)}$, $\hat{\bar{c}}_a^{Direct*(b)}$ and $\hat{\bar{c}}_a^{noZ*(b)}$.

6. Calculate the Monte-Carlo approximation to the theoretical bootstrap estimator for each area

$$MSE^*(\hat{\bar{c}}^{EB}) = \frac{1}{B}\sum_{b=1}^{B}\left(\hat{\bar{c}}_a^{EB*(b)} - \bar{c}_a^{*(b)}\right) \tag{2.13}$$

with analogous measures for the other estimators.

While models omitting information on $\mathbf{Z}_a$ by necessity perform worse than EB, we use this comparison to show the relative importance of including credit card data and energy use in improving the accuracy of our small-area estimates.

2.3 Standard errors

The 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 the model in 2.7 using maximum likelihood in our consumption sample. Obtain estimators of $\boldsymbol{\beta}$,$\boldsymbol{\pi}$,$\mathbf{u}_a$, σ_ε^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 this to construct

$$\log c_{i,a}^{\dagger(b)} = \mathbf{X}_{i,a}\hat{\boldsymbol{\beta}} + \mathbf{Z}_a\hat{\boldsymbol{\pi}} + \hat{u}_a + \varepsilon_{i,d}^{\dagger(b)} \tag{2.14}$$

4. For each bootstrap population, take values from with the same indices as the initial sample and use these to produce the EB estimators as described in Section 2.1, $\hat{\bar{c}}_a^{EB\dagger(b)}$.

The standard error associated with each area level estimate is then

$$SE(\hat{\bar{c}}_a^{EB}) = \frac{1}{B}\sum_{b=1}^{B}\left(\hat{\bar{c}}_a^{EB\dagger(b)} - \hat{\bar{c}}_a^{EB}\right)^2 \tag{2.15}$$

The standard errors for the other estimators are computed in an analogous way (replacing the model in 2.14 with the appropriate estimation model).

3 Data

Living Costs and Food Survey: Our source of consumption data is the LCFS, which is an annual survey of the spending patterns of households in Great Britain. 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 10,243 households across the two years.

The LCFS contains aggregated COICOP category level expenditures which are derived from the spending diary and recall questions. Our main outcome variable of interest is total expenditure, which sums across all COICOP groups. However, we also consider total expenditure less housing costs (which we obtain by excluding spending on rent and mortgage interest). 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: We use geocoded-versions of the 2018 and 2019 Annual Population Survey (APS) as our auxilliary dataset. The Annual Population Survey is a continuous household survey, conducted with the aim of providing information on social and labour market variables at a local level in the years between the (decennial) UK censuses. We apply the same restrictions to households and derive the same explanatory variables as in the LCFS. (The APS does not contain consumptions measures.) The final sample size after cleaning is much larger than the LCFS - around 387,000 households over the two years. We pool across two calendar years to increase the precision of our estimates.

Card transaction data and energy-use: We complement the household-level covariates from the APS with area-level covariates derived from two other data

sources. The first source of data is a transaction-level credit and debit card dataset provided by Fable Data. The data contains on average transactions associated with approximately 1.9 million credit card accounts and 201,000 debit card accounts in each year. Across the 2018 and 2019 calendar years, our transaction data comprises £4.9 billion of expenditure. In addition, the data contains demographic information, including the postal sector of the account holder's home address. In order to utilise these in the EB framework, we construct local authority mean debit and credit card expenditure, and use these as area-level covariates in the nested error model 2.7. In addition to the APS and Fable data, we also use data on average energy consumption at the outcode level (the first 2-4 digits of the postcode), published by the Department for Energy, Security and Net Zero. We include outcode-level mean electricity consumption (in KWH) as another area-level control in the model.

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 distribution of consumption and income (Office for National Statistics, 2024). The GDHI is advantageous 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 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 the model outlined in Section 2 must be available 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 equivalise our estimates of average consumption for family size, according to the OECD equivalence scale. The Z variables we use are average credit and debit card spending in the households' LAD, and average electricity consumption in KWh for the households' local postcode.

3.2 Linear regression synthetic estimator

To assess the importance of area-specific residual information in estimating local consumption, we also estimate an analogue of 2.7 excluding the random effect. Specifically, we can estimate the following model:

$$\begin{aligned} \log c_{i,a} &= \mathbf{X}_{i,a}\boldsymbol{\theta} + \mathbf{Z}_a\boldsymbol{\delta} + \nu_{i,a} \\ \nu_{i,a} &\sim_{iid} N(0, \sigma_\nu^2) \end{aligned} \tag{3.1}$$

We estimate this model using OLS and use the coefficient estimates to generate imputed consumption for each household in the APS. As with the EB procedure, we use a Monte-Carlo procedure to transform our imputed values of log consumption in the APS into levels before taking means. The estimates produced by this model will be identical to those estimates by EB in areas where the LCFS sample size is zero.

We report the results of a random-effects regression model of household total consumption spending in the LCFS on these variables, estimated using maximum-likelihood, in Table A.1 in the Appendix.[4] We also include results of an F-test comparing model 2.7 to model 3.1. This test has a p-value of < 0.001 meaning we strongly reject the null of no area-level random effects.

[4]We also report the parameter estimates for the regression excluding housing-related expenditures in Table A.2.

4 Estimates of local consumption spending

Figure 4.1 shows our EB estimates of weekly equivalised consumption for local authorities in Great Britain, generated from a sample of individuals interviewed between 2018 and 2019. We find that average consumption estimates range between £216 per week (in Leicester) and £512 (in City of London). The median average level of equivalised weekly consumption is £307. Areas in the South East of England, such as Windsor and Maidenhead (£436), have relatively high levels of equivalised consumption, in contrast to areas such as Sandwell (£228) and Dudley (£238) in the West Midlands. Within London, we observe significant variation across local authorities. Barking and Dagenham (£236) exhibits the third lowest level of consumption nationally, whilst residents of Richmond upon Thames (£403) have the seventh highest level of consumption (and second highest in London).

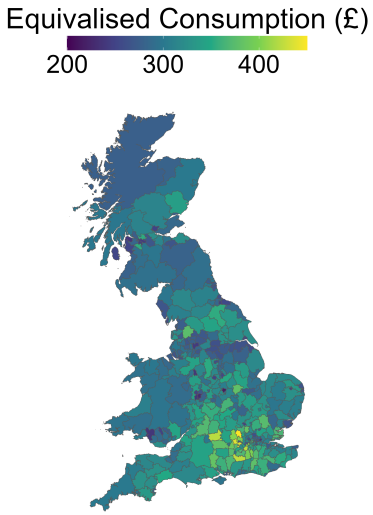


Figure 4.1: *EB estimates of 2018/2019 Equivalised Consumption*

Notes: To facilitate visual comparisons, we truncate the distribution of equivalised consumption so that the City of London has a consumption value equal to £450 (equal to the second highest-consuming local authority).

We also estimate equivalised weekly consumption excluding expenditure on rent and mortgage interest to give an 'after housing costs' measure of consumption spending. We plot these estimates on a map in Figure B.1 in the Appendix. Figure 4.2 plots estimated consumption including and excluding expenditure on housing-associated expenditures against each other, with a linear fit plotted through the observations. The five local authorities with the largest fall in consumption after excluding housing-related costs are all in London. These differences likely reflect higher housing costs in these localities.

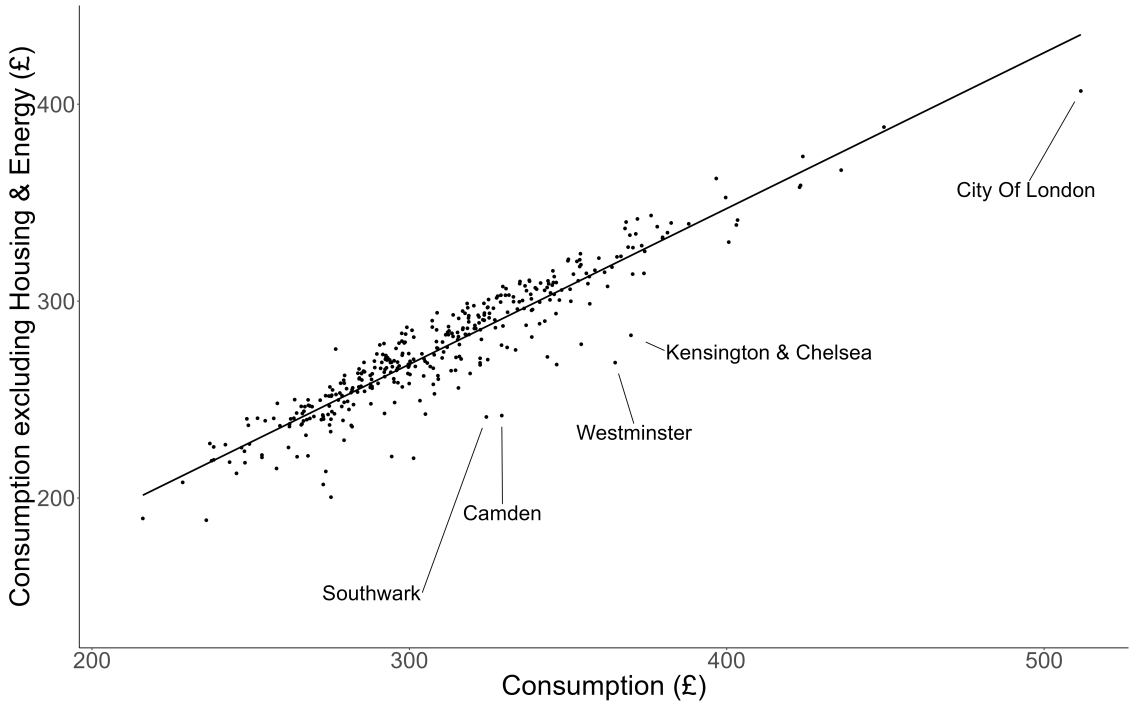


Figure 4.2: *EB estimates: including and excluding housing-related costs*

As outlined in Section 2.3, we obtain area-specific standard errors for our EB estimates, generated via a parametric bootstrap procedure. In Figure 4.3, we plot our estimates of equivalised consumption (including housing costs) with the corresponding 95% confidence intervals, ordered by rank in the equivalised consumption distribution. 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 4.1, however, the EB estimator substantially improves upon the sample mean estimator in terms of variance.

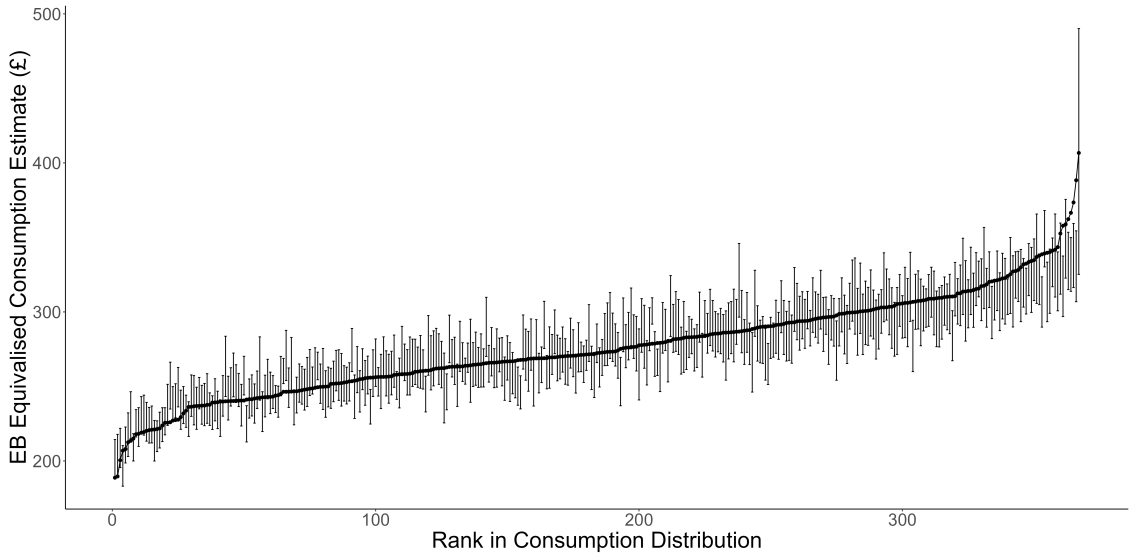


Figure 4.3: *EB estimates of 2018/2019 Equivalised Consumption with 95% confidence bands*

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

4.1 Comparison of estimation methods

We use the 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 the credit and debit card transaction data improves the estimator's accuracy. Figure 4.4 plots the unweighted average (across local authorities) of these mean squared errors across areas, by estimation method. The average MSE of the EB estimator is substantially lower than that of the direct sample mean estimator.[5] Including the credit and debit card data as covariates in the EB estimation method reduces the mean squared error by 22.0%. This is consistent with our finding (reported in Table A.1 in the Appendix) that credit-card expenditure is a highly significant predictor of consumption.

In Figure 4.5, we plot area-specific mean squared errors corresponding to the sample mean and EB estimator (including the credit card data), against the LCFS sample size. Intuitively, for both 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

[5] All 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.

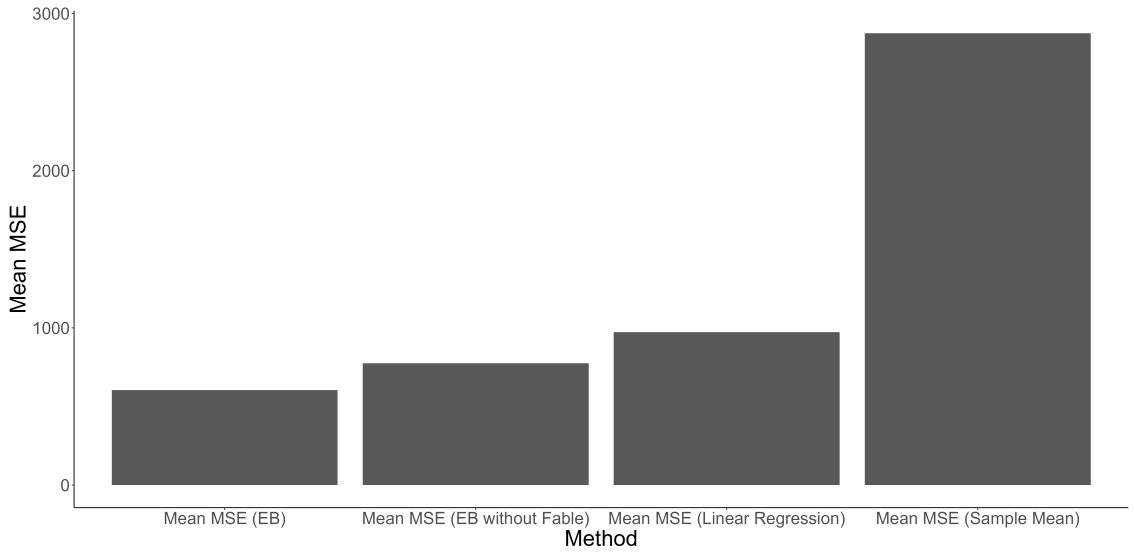


Figure 4.4: *Average Mean Squared Error by Estimation Method*

depicted in Figure 4.4 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).

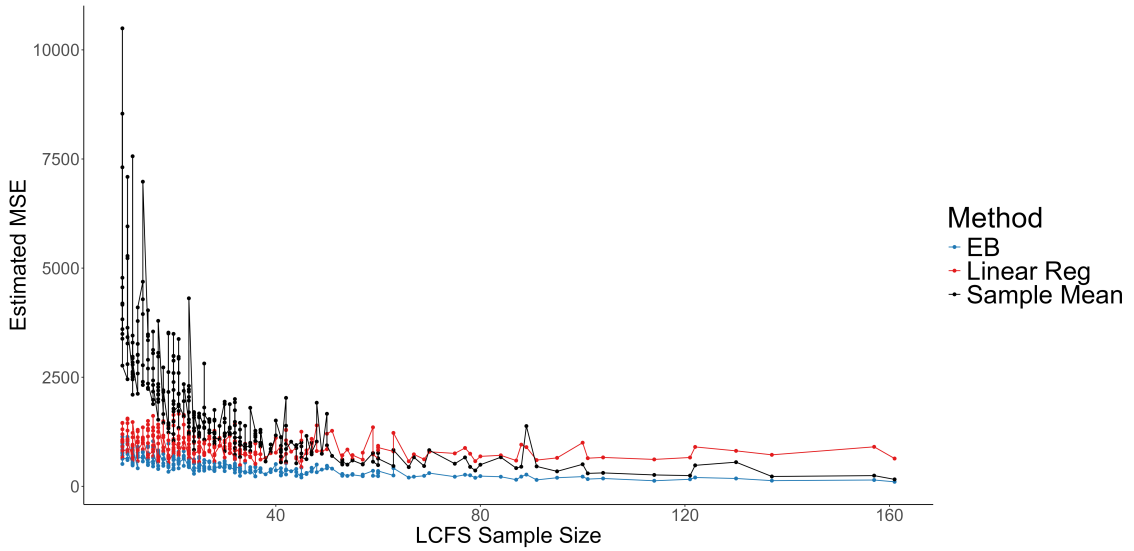


Figure 4.5: *Area-Level MSE, by Method*

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 2.11 as:

$$MSE(\hat{\bar{c}}_a^{EB}) = V(\hat{\bar{c}}_a^{EB} - \bar{c}_a) + \{E(\hat{\bar{c}}_a^{EB} - \bar{c}_a)\}^2 \tag{4.1}$$

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

$$MSE(\hat{\bar{c}}_a) = V(\hat{\bar{c}}_a^{Direct} - \bar{c}_a) + \{E(\hat{\bar{c}}_a^{Direct} - \bar{c}_a)\}^2 \tag{4.2}$$

Note that because the target parameter $\bar{c}_a$ is a random variable (varying across draws from the superpopulation model 2.7), 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 model bias and the variance of the *estimation errors*. We plot this decomposition in Figure 4.6. 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 imprecisely estimated. The linear regression estimator exhibits larger absolute biases than the EB estimator for 240 out of 305 reported areas, and the sample mean is more biased than the EB estimator for all but 18 local authorities. As the sample size grows, the bias of the sample mean approaches zero, as expected.

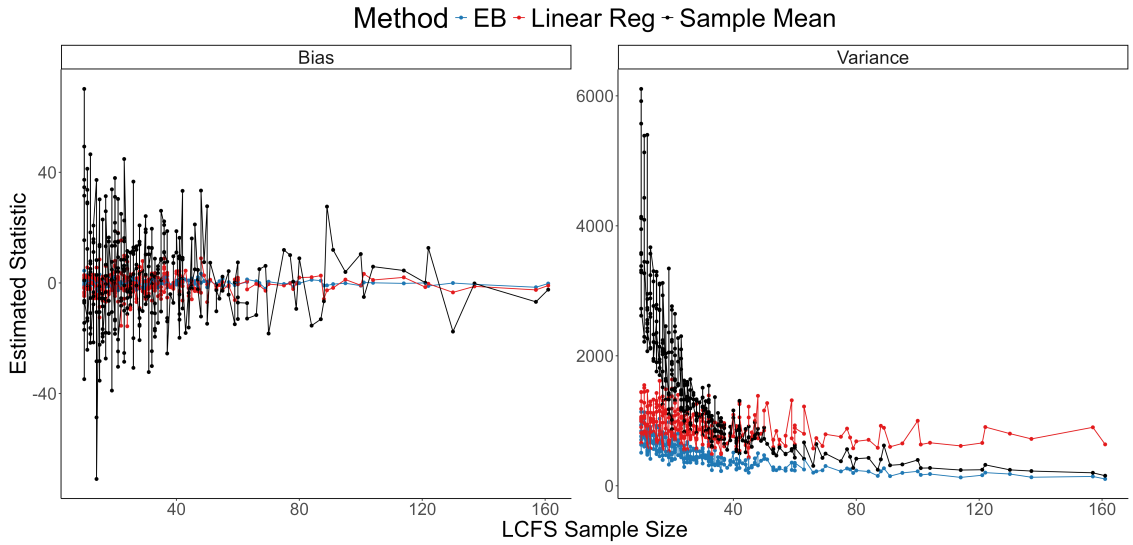


Figure 4.6: *Area-Level Decomposition of MSE, by Method*

4.2 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)), and outline this procedure in Section 3.

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 be be more likely to capture very high income households that might be less likely to respond to the consumption survey than 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 affecting as great an effect on our consumption measure.[6]

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 may

[6]In Figures B.3 and B.4 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. In line with this, our consumption ranks are more strongly correlated with the Income IMD ranking than the GDHI income ranking, as shown by the correlation matrix in Table B.2.

differ 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 consumption against the mean per-capita income distribution, along with a linear fit, in Figure 4.7. Of particular note is the disparity between the position of London areas in the income and consumption distributions. For example, Brent ranks at the 81st percentile of the local authority distribution of income per capita, but at the 6th percentile of the corresponding equivalised consumption distribution. Of the ten areas with the largest absolute difference between their ranks on measures of income and consumption, nine are located in London. We also compare the relative rankings of areas in the net-of-housing consumption distribution and the income distribution in Figure B.2 in the Appendix. Since the GDHI is measured before rental and mortgage principal payments, we are unable to compute net-of housing income.

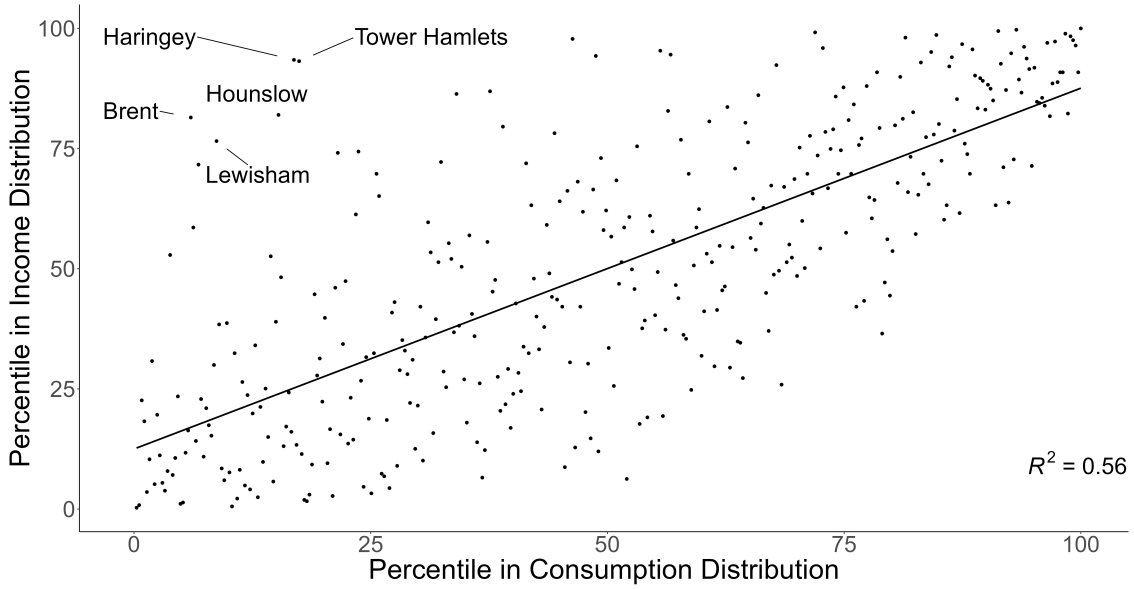


Figure 4.7: *Equivalised Consumption vs per-capita income - ranks*

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 consumption is 35%. This qualitative pattern holds for all the ratios reported in Table 4.1, and is starkest at the extremes: individuals in the richest local authority (City of London) have incomes 16.2 times higher than those in the poorest local authority (Leicester), but only consume 2.4

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. Nonetheless, 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	Equivalised Mean Consumption	2.35	1.47	1.35	1.18
2	Mean Income per Capita	16.18	1.84	1.65	1.27

Table 4.1: *Percentile Ratios of local authorities*

5 Conclusion

This paper 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.

Transaction 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 full short of giving a complete picture of households' consumption levels. Our findings suggest that transaction 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.

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.

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 (2020). Development of regional household expenditure measures - Office for National Statistics. Technical report, Office for National Statistics.

Office for National Statistics (2024). Regional gross disposable household income, UK: 1997 to 2022. Technical report.

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.

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: Total Consumption*

Variable	Estimate	Std. Error	T value	P value
(Intercept)	0.999	0.53	1.886	0.059
HH contains an U14 individual	0.047	0.027	1.72	0.085
HH contains an individual aged 14-18	0.089	0.058	1.523	0.128
HH contains >1 adult	0.44	0.019	22.579	<0.001
No. U14s in household	0.015	0.014	1.043	0.297
No. 14-18 in household	0.082	0.047	1.766	0.077
No. additional adults	0.204	0.012	16.77	<0.001
Sex of HoH	0.009	0.012	0.744	0.457
Log LAD-average credit card spending	0.515	0.096	5.362	<0.001
Log LAD-average debit card spending	0.081	0.05	1.613	0.107
HoH finished education between 18 and 21	0.161	0.017	9.403	<0.001
HoH finished education aged over 21	0.219	0.017	12.917	<0.001
HoH is still in education	0.286	0.051	5.64	<0.001
HoH ethnicity is White	0.098	0.038	2.615	0.009
HoH ethnicity is Asian/Asian British	-0.158	0.046	-3.446	0.001
HoH ethnicity is Black African/Caribbean/Black British	-0.15	0.053	-2.848	0.004
HoH is aged 30-40	0.037	0.025	1.494	0.135
HoH is aged 40-50	0.012	0.025	0.476	0.634
HoH is aged 50-60	-0.012	0.025	-0.46	0.645
HoH is aged 60-70	-0.061	0.026	-2.313	0.021
HoH is aged 70+	-0.257	0.027	-9.38	<0.001
HH are mortgage holders	0.149	0.015	10.176	<0.001
HH are owner occupiers	0.181	0.016	11.594	<0.001
LAD-median electricity consumption (MWH)	0.091	0.018	5.006	<0.001
2018 Year Dummy	-0.048	0.014	-3.447	0.001
F-test for Random Effect	70.552			<0.001
σ^2_ϵ	0.288			
σ^2_μ	0.008			
Marginal R^2	0.392			
Conditional R^2	0.409			

Notes: These parameter estimates correspond to a random effects regression of log consumption on covariates and a local-authority random effect. It was estimated on a sample of 10243 households from the 2018 and 2019 Living Costs and Food Survey. The F-test reports a Chi-Squared statistic testing the random effects model against a linear model omitting the random area effect. The marginal R^2 denotes the share of the variance explained by the (fixed) covariates, and the conditional R^2 denotes the share of the variance jointly explained by the fixed and random effects.

Table A.2: *Regression Results: Consumption excluding housing-related expenditures*

Variable	Estimate	Std. Error	T Value	P Value
(Intercept)	1.524	0.526	2.897	0.004
HH contains an U14 individual	0.031	0.029	1.05	0.294
HH contains an individual aged 14-18	0.118	0.062	1.898	0.058
HH contains >1 adult	0.447	0.021	21.648	<0.001
No. U14s in household	0.047	0.015	3.138	0.002
No. 14-18 in household	0.09	0.049	1.814	0.07
No. additional adults	0.224	0.013	17.353	<0.001
Sex of HoH	-0.011	0.012	-0.93	0.352
Log LAD-average credit card spending	0.389	0.096	4.065	<0.001
Log LAD-average debit card spending	0.064	0.05	1.288	0.198
HoH finished education between 18 and 21	0.133	0.018	7.344	<0.001
HoH finished education aged over 21	0.189	0.018	10.521	<0.001
HoH is still in education	0.216	0.054	4.018	<0.001
HoH ethnicity is White	0.159	0.04	3.988	<0.001
HoH ethnicity is Asian/Asian British	-0.202	0.049	-4.148	<0.001
HoH ethnicity is Black African/Caribbean/Black British	-0.135	0.056	-2.421	0.015
HoH is aged 30-40	0.052	0.026	1.998	0.046
HoH is aged 40-50	0.054	0.026	2.04	0.041
HoH is aged 50-60	0.067	0.027	2.53	0.011
HoH is aged 60-70	0.037	0.028	1.319	0.187
HoH is aged 70+	-0.155	0.029	-5.33	<0.001
HH are mortgage holders	0.293	0.015	18.908	<0.001
HH are owner occupiers	0.453	0.017	27.333	<0.001
LAD-median electricity consumption (MWH)	0.091	0.019	4.85	<0.001
2018 Year Dummy	-0.029	0.014	-2.052	0.04
σ^2_ϵ	0.325			
σ^2_μ	0.006			
Marginal R^2	0.405			
Conditional R^2	0.417			

Notes: These parameter estimates correspond to a random effects regression of log consumption excluding housing-related expenditures on covariates and a local-authority random effect. It was estimated on a sample of 10243 households from the 2018 and 2019 Living Costs and Food Survey. We exclude rent and mortgage interest payments from consumption in line with the HBAI definition of housing costs. The marginal R^2 denotes the share of the variance explained by the (fixed) covariates, and the conditional R^2 denotes the share of the variance jointly explained by the fixed and random effects.

Appendix B Additional Tables and Figures

Table B.1: *Covariate Balance between LCFS and APS samples*

Variable	LCFS Sample Mean	APS Sample Mean
No. U14s in household	0.409	0.41
No. 14-18 in household	0.097	0.108
No. additional adults	0.816	0.844
Sex of HoH (1 = male)	0.596	0.589
HH contains an U14 individual	0.246	0.243
HH contains an individual aged 14-18	0.083	0.091
HH contains >1 adult	0.673	0.666
LAD-average credit card spending	835.048	820.745
LAD-average debit card spending	1012.545	996.36
HoH ethnicity is White	0.92	0.916
HoH ethnicity is Asian/Asian British	0.039	0.041
HoH ethnicity is Black African/Caribbean/Black British	0.02	0.022
HoH ethnicity is Mixed Multiple Ethnic Groups/Other Ethnic Group	0.022	0.021
HoH is aged 18-30	0.072	0.068
HoH is aged 30-40	0.159	0.153
HoH is aged 40-50	0.173	0.17
HoH is aged 50-60	0.194	0.197
HoH is aged 60-70	0.182	0.18
HoH is aged 70+	0.221	0.231
HH are mortgage holders	0.301	0.296
HH are owner occupiers	0.374	0.38
HH are renters	0.325	0.324
LAD-median electricity consumption (MWH)	2.942	2.899
HoH finished education before 18	0.574	0.593
HoH finished education between 18 and 21	0.196	0.198
HoH finished education aged over 21	0.212	0.201
HoH is still in education	0.018	0.008

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

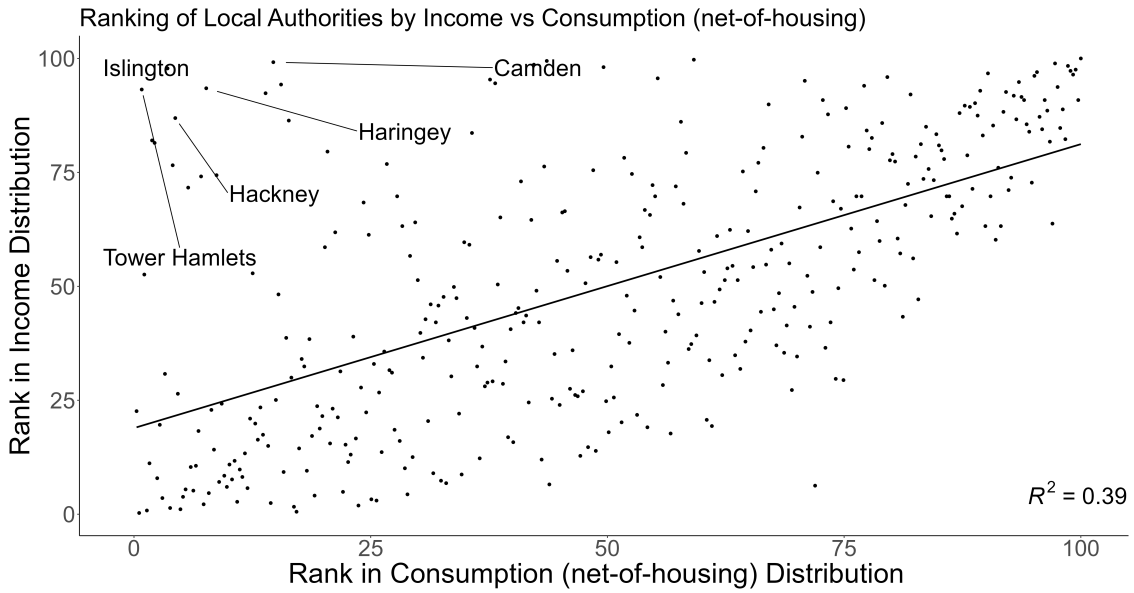


Figure B.2: *Equivalised Consumption (net of housing costs) vs per-capita income - ranks*

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.

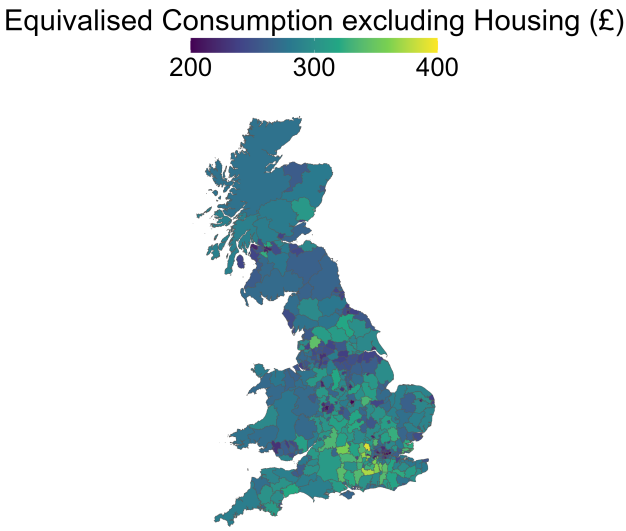


Figure B.1: *EB estimates of 2018/2019 Equivalised Consumption - excluding housing-related expenditures*

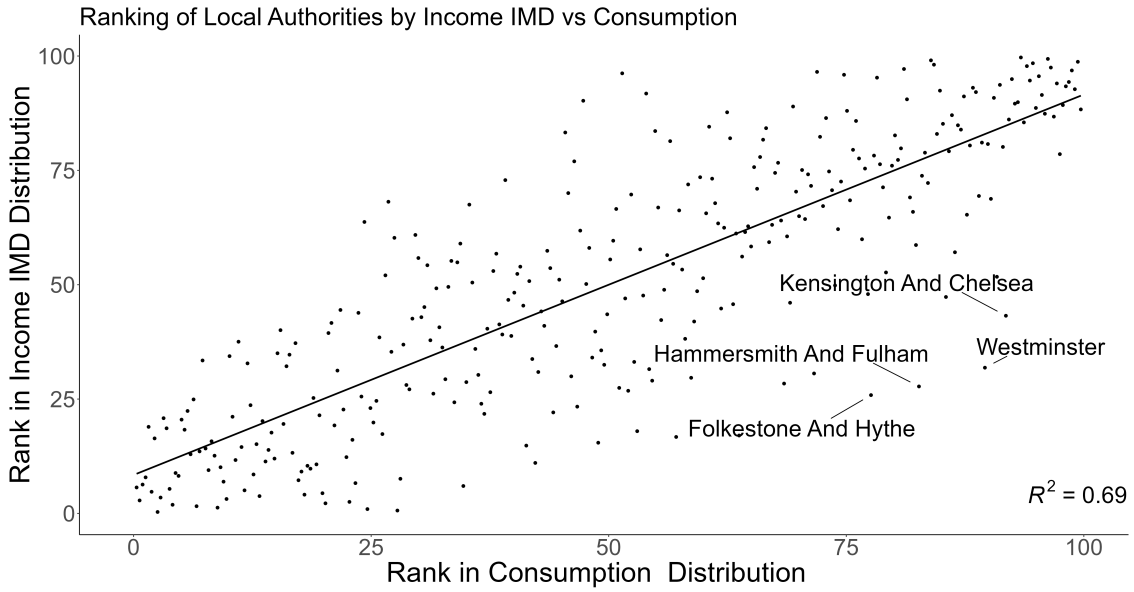


Figure B.3: *Equivalised Consumption vs Income IMD - ranks*

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.

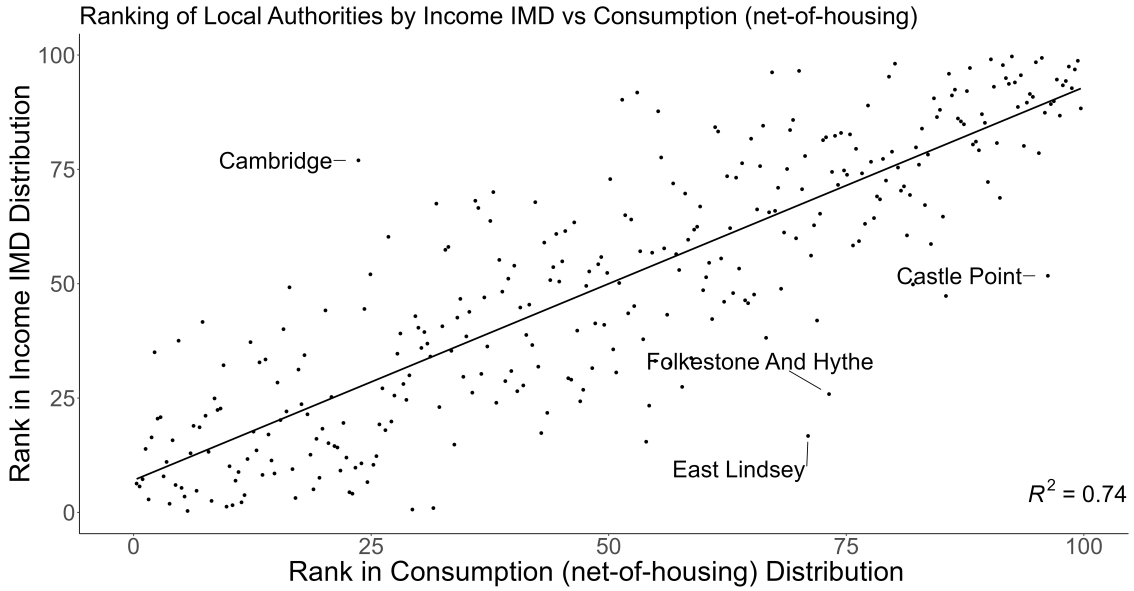


Figure B.4: *Equivalised Consumption (net of housing costs) vs Income IMD - ranks*

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.

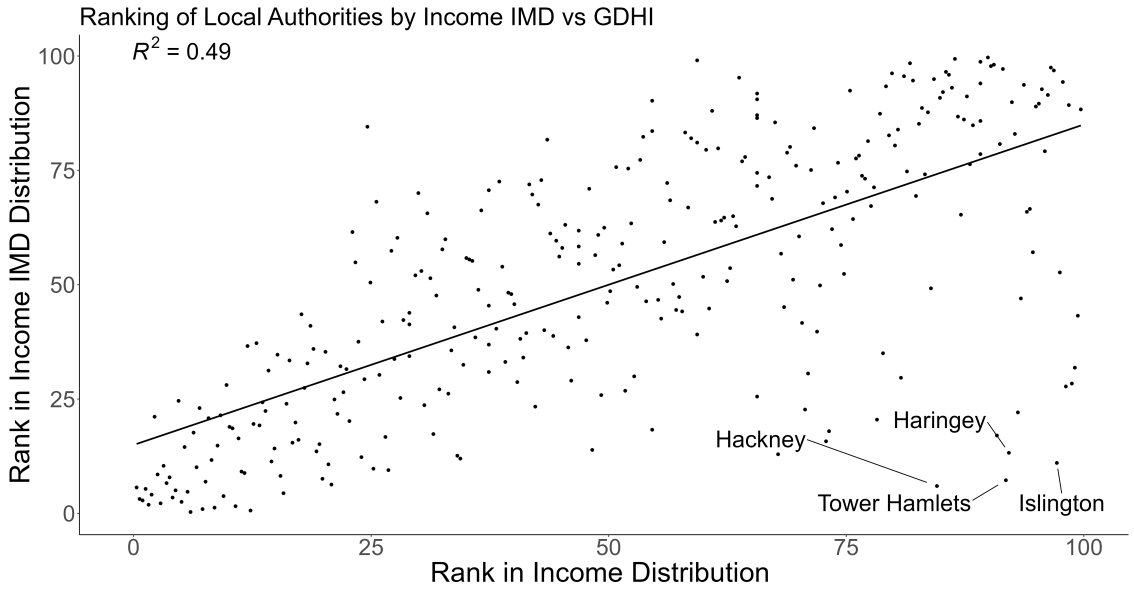


Figure B.5: *GDHI Income vs Income IMD - ranks*

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.

	Consumption	Net Consumption	Income IMD	GDHI
Consumption	1.000	0.927	0.832	0.757
Net Consumption	0.927	1.000	0.859	0.627
Income IMD	0.832	0.859	1.000	0.701
GDHI	0.757	0.627	0.701	1.000

Table B.2: *Correlation Matrix of local authority ranks by consumption, Income IMD and GDHI Income*

Notes: 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. We label the four local authorities with the largest absolute disparity in rank between the two measures. Net Consumption refers to our estimates of consumption after housing expenditures are removed.

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