Spatial microsimulation estimates of household income distributions in London boroughs, 2001 and 2011

Alex Fenton

Contents

List of Tables

List of Figures

Centre for Analysis of Social Exclusion

The Centre for Analysis of Social Exclusion (CASE) is a multi-disciplinary research centre based at the London School of Economics and Political Science (LSE), within the Suntory and Toyota International Centres for Economics and Related Disciplines (STICERD). Our focus is on exploration of different dimensions of social disadvantage, particularly from longitudinal and neighbourhood perspectives, and examination of the impact of public policy.

In addition to our discussion paper series (CASEpapers), we produce occasional summaries of our research in CASEbriefs, and reports from various conferences and activities in CASEreports. All these publications are available to download free from our website. Limited printed copies are available on request.

For further information on the work of the Centre, please contact the Centre Manager, Jane Dickson, on:

Telephone: UK+20 7955 6679
Fax: UK+20 7955 6951
Email: j.dickson@lse.ac.uk
Web site: http://sticerd.lse.ac.uk/case

Editorial note and acknowledgments

Abstract

Spatial microsimulation (SMS) is a range of techniques for estimating the local distribution of a variable – here, household income – by combining social survey microdata with Census or administrative population totals. This paper makes a case for the value of these methods in social policy analysis of spatial economic differences because unlike other methods and sources, they permit distributional analysis of income, encompass both market outcomes and secondary distribution through taxes and transfers, and measure income poverty in standard national terms. As a demonstration of spatial microsimulation by iterative proportional fitting (IPF), the household income distribution in London's 33 boroughs in 2001/02 and 2011/12 is estimated in this paper. The coherence and plausibility of the results in comparison to other official statistics is examined in some detail. Two refinements to standard IPF methods are presented, including "multilevel IPF", which allows the use of both person- and household-level data; this is found to improve the estimation of poverty rates. The paper confirms the value of SMS for synchronic spatial analysis, and argues for its hitherto little-explored use in modelling spatial differences in the effects of fiscal and welfare policy changes.

Key words: microsimulation, income estimation, poverty, spatial inequality

JEL classification: C81

Corresponding author: alex.fenton@pressure.to

Glossary of Acronyms

ASHE Annual Survey of Hours and Earnings: UK survey source on employee wages and salaries

DCLG Department for Communities and Local Government

DWP Department for Work and Pensions

GDHI Gross Disposable Household Income: national accounting measure of average household income

FRS Family Resources Survey: UK government survey of household income

HBAI Households Below Average Incomes: UK survey of the income distribution, derived from FRS

HHO here, simulation using household data only (see section 3.2)

HMT HM Treasury

MLV here, multi-level simulation using household and individual data (see section 3.3)

IFS Institute for Fiscal Studies

IPF Iterative Proportional Fitting

LFS Labour Force Survey

NUTS2 / NUTS3 Nomenclature of Territorial Units for Statistics areas: standardised European statistical regions

NS-SEC National Statistics Socio-Economic Classification

RMSE Root Mean Squared Error

SMS Spatial Microsimulation

SPI Survey of Personal Incomes: UK statistical source on individual incomes from tax records; here, simulation using this source (see section 3.4)

UMBR Unadjusted Means-Tested Benefits Rate: poverty proxy measure using means-tested social security benefits

1 Introduction

One feature of Britain's economic and social landscape is the existence of deep, persistent, and even growing spatial disparities in income and wealth between regions, cities and neighbourhoods (Dorling et al. 2007; Gardiner et al. 2013). These disparities have been a recurring, if not constant, presence both in government policy and in research for much of the last century (O'Hara 2005). Two broad disciplinary and practical traditions might crudely be distinguished within the field: *regional economics*, and the *geography of deprivation*. These traditions have their distinct vocabularies, methods and data sources, and have not always been in communication with one another.

An important strand of *regional economics* describes, analyses and explains disparities between regions or between smaller spatial units. This includes analyses of differentials in output, wages or employment rates. The measures of economic well-being used in this work are, however, typically aggregates; they provide a single figure for each geographical unit of interest at each point in time. Seen through the lens of social policy, these approaches have the shortcoming that they do not consider the *distribution* of the aggregate within each area (T. Atkinson, Guio, and Marlier 2015). An increase in Gross Disposable Household Income (GDHI) in an area need not correspond to any reduction in the proportion of households in relative income poverty, if the increase is so distributed that it does not improve the position of those below and around the poverty line. Even analysis of wages, where sources such as the Annual Survey of Hours and Earnings (ASHE) permit some investigation of the distribution, considers only at the primary distribution of income via the labour market. Such analysis can take account of neither the effects of the secondary distribution of income through taxes and transfers nor of housing costs on the final living standards of households. A policy which, through income transfers, increased the proportion of households achieving a minimum consumption standard would, for example, not be reflected in wages data.

In the *geography of deprivation* the focus has typically been upon smaller spatial units, often neighbourhoods. The characterisation and measurement of poverty and deprivation is the primary focus, rather than the relative wellbeing of the whole resident population. Key quantitative tools for such research are a variety of deprivation indices dating back to the 1970s. Earlier incarnations concentrated on the incidence of material deprivation, and used Census data; from the mid-1990s, indices have sought to characterise other domains of exclusion, such as sickness and crime, often through the addition of administrative data. Such indices have become widely used in policy circles and applied research, and can be adequate and useful proxies for the relative distribution of income deprivation (Fenton 2013b). They do, however, also detach small-area analysis of deprivation from the analysis of the whole income distribution, and, further, from the terms, techniques and measures in which low-income and inequality are addressed at national level. In addition, as a result of both welfare reform and rises in in-work poverty, the social security administrative data upon which present indices depend increasingly weakly identify the overall prevalence of low income (Fenton 2013a). This argument is acknowledged (and then dismissed), for example, in the government's re-

sponse to the consulation on updating the English Indices of Deprivation (Department for Communities and Local Government 2015, pp.18-19).

Spatial microsimulation (SMS) is a potential solution to at least some of these problems. SMS involves the synthesis of microdata from national household income surveys and local area population totals from censuses or administrative data. The outputs of SMS are effectively new datasets, which can be used to estimate income in the terms used in the specialised social surveys, but at the spatial level of districts or small-areas rather than national or regional data originally available. A wide variety of descriptive statistics can be calculated by SMS – not only means, but quantiles, poverty rates using thresholds, and measures of dispersion and inequality. Since the survey data separately measure income components, such as wages and transfers, SMS estimates likewise reflect both the primary and secondary distributions of income on household economic circumstances, and are sensitive to changes in transfers and taxation rates. The effects of housing costs on final income consumption are similarly incorporated. Simulation results also permit a degree of sub-group analysis, such as, for example, the separation of in-work and out-of-work poverty, or the incidence of child poverty. For social policy analysis of poverty, SMS addresses criticisms of aggregate regional measures, that they do not have any necessary correlation with poverty, and of deprivation indices, that they poorly reflect in-work poverty.

Non-spatial simulations with survey data are used in Britain in both the Treasury (HMT) and the Department for Work and Pensions (DWP), as well as in research institutes like the Institute for Fiscal Studies (IFS) (S. Edwards and Gault 2009; Roe and Rendle 2009). These simulations model the effects of policy changes on households. *Spatial* microsimulation techniques have, however, not thus far found widespread application in official statistics and analysis. The Department for Communities and Local Government (DCLG) commissioned a pilot project to examine using simulated income estimates within the English deprivation indices; this resulted in a very useful working paper (Anderson 2007), but the actual results and resulting datasets were not made publicly available. Only in Northern Ireland are the results in the public domain: (Anderson 2008). The Office for National Statistics' *Small Area Income Estimates* are produced by econometric methods, discussed briefly below, which limit both the flexibility of the outputs and the reproduction and adaptation of the methods outside government. Since the decennial Census has not and likely will not measure income (Office for National Statistics 2005), there remains a want of detailed and readily available sub-regional estimates of household income.

An exciting further possibility for policy research, which is however only lightly treated here, is using SMS results to simulate the local effects of national policy changes to taxes or benefits, effectively giving spatial disaggregation to the kind of modelling done by DWP, HMT and IFS. One could produce a simulation which presented to MPs the precise effects of proposed tax increases or benefit cuts in their constituency![1] Finally, there may also be non-scientific reasons for microsimulation to enjoy a conjuncture.

---

[1] I am grateful to Ben Anderson for this suggestion.

Microsimulation combines existing datasets rather than the demanding expensive new fieldwork, thus offering the promise of getting "more for less (Haslett et al. 2010), in a period since 2010 in British official statistics where there has been a marked tendency to cut, rather than extend, primary fieldwork in social surveys. Britain may also be propitious territory for microsimulation given the reliance of the UK, as other anglo-saxon statistical systems, on social surveys (Snorrason, Byfuglien, and Vihavainen 2015). The availability of relatively rich survey data and the absence (unlike in many nordic and continental European statistical systems) of population register data is well suited to the microsimulation approach.

1.1 Outline and purpose of the paper

This paper presents a worked spatial microsimulation of household incomes for London boroughs in 2001/02 and 2011/12, with the hope that it might pique the interest of social policy researchers in the method. The paper considers a number of modest technical adaptations to the IPF fitting technique, which may be of interest to those already familiar with SMS. The paper's central focus is on the plausibility of the SMS results as estimates of the whole income distribution and lower and upper incomes. London is here a useful example, since there are found there high rates of poverty, marked inter-area differences, and the highest incomes in Britain. The differences between boroughs and over time are treated here, however, from a largely abstract and technical perspective. A companion paper, *Gentrification in London: a Progress Report 2001-2013*, takes up the interpretation of the results in spatial and policy context, and it will be more relevant for those with a primarily substantive interest in income and poverty in London.

The body of the paper first sets out more formally what SMS is, and how it differs from other small-area estimation approaches. The sources and techniques used are then described, including two refinements to microsimulation techniques which are new in the published UK literature. The largest part of the paper is an examination of the coherence of the results with other available sources, including published poverty measures, wages and benefits data.

2 Spatial microsimulation and other methods

Spatial microsimulation combines social survey data with local Census or administrative data. The survey data – in this example, the *Family Resources Survey* (FRS) and the derived dataset *Households Below Average Income* (HBAI) – include the variables whose local distribution one ultimately wishes to estimate. In our case, these are the various detailed measures of gross, net and equivalised income contained in FRS. The local data are population totals (or "margins") for each area of interest, here 33 London boroughs, for a set of household and person characteristics that predict income, such as employment status, housing tenure and occupational class. The local data and survey cases are connected by a set of *link variables*, which are measured in the same way

in both datasets: if the survey dataset measures employment status for each case using three categories, employed, unemployed and inactive, then the local data must provide area population totals for these same categories, identically derived.

With these two datasets to hand, the survey data are fitted, or constrained, to the local population margins in each borough successively. The weights of each survey case are adjusted until the weighted survey population resembles, as near as possible, the characteristics of the borough. Several distinct techniques exist to do this re-weighting of the survey cases to the local population totals, including generalised regression weighting and combinatorial optimisation (a useful summary is provided by Whitworth 2013, p.5ff). Here iterative proportional fitting (IPF) is used (Anderson 2012). In IPF, the weights of each case are scaled arithmetically to each set of marginal totals constraint in turn, and this process then repeated a desired of times. This re-weighting is fundamentally the same process as that of 'calibration' in the literature of survey methodology, from where it was borrowed, and where IPF specifically is often termed 'raking' (Lumley 2004). IPF is an arithmetically simple technique known since the 1950s, and is well described in the literature, with step-by-step illustrations; further elaboration is thus here foregone (Ballas et al. 2005; Simpson and Tranmer 2005).

The core of the microsimulation approach is thus the *linking variables* which are, firstly, available in the survey data cases and as population totals for the areas of interest and, secondly, are predictors of the outcome variable of interest (income) in the survey source. Once the IPF has been carried out, one has a set of new weights, one for each survey case in each local area of interest. These new survey weights are then applied to the survey data to provide (weighted) estimates of the income and poverty statistics of interest for each local area.

2.1 *Other small-area income estimation methods*

It is worth noting in passing that microsimulation is not the only method by which local income estimates can be derived through synthesis of local area and survey data. In recent years several reviews have compared microsimulation to other small-area estimation methods from both more theoretical and more practical standpoints (Useful reviews include Betti and Lemmi 2013; New Zealand and Statistics New Zealand 2010; Whitworth 2013). This literature typically distinguishes a family of methods under the heading *spatial microsimulation* from a set of approaches labelled *modelled* or *statistical*.

These have in common some statistical method for combining the best available survey measure of income – in the UK, normally the *Family Resources Survey* or *Understanding Society* – with local area characteristics which predict inter-area variation in that income. Modelled or statistical approaches are so called because a regression model is first fitted to the sample survey which estimates how some statistic of income, such as its mean, is related to local characteristics. Once a model has been fitted, the param-

eters are used to give predicted values for all areas from the local data. UK examples include Bramley's work on housing affordability (Bramley and Lancaster 1998; Wilcox and Bramley 2010) and the Office of National Statistics' poverty estimates for small areas (Methodology Directorate and Centre for Regional and Local Statistics 2010), along with ONS's comparable earlier estimates of mean income. Parametric approaches – that is, those based on a formal specification of the statistical income distribution and estimation of its parameters - are also widely used in development economics (a key paper here is Elbers, J. O. Lanjouw, and P. Lanjouw 2003).

Aside from *modelled* and *simulated* approaches, one might also note the in applied statistics of *heuristic* approaches, in which two or more data sources are synthesised or calibrated to produce local estimates, but with a technique not formally based on statistical theory. Such approaches apply judgement to the synthesis of multiple sources, translating, for example, observed variance in one dataset to predicted variance in another. The Greater London Assembly's suite of small-area income estimates nicely exemplify this approach (GLA Intelligence 2014).

To conclude the digression: there is not at present any decisive statistical criterion for choosing between modelling and simulation[2]. In any case, modelled and simulated approaches have been argued to share an underlying model of the relationships among variables (Haslett et al. 2010). Preferences for one method or another are partly disciplinary: researchers with a statistical science background tending to estimations based on distributions, econometrically oriented researchers preferring methods employing a predictive model of incomes in micro-data (e.g. Methodology Directorate and Centre for Regional and Local Statistics 2010), and so on. The paper here takes a narrative and exemplifying approach to the method and concentrates not on the underlying model, but on the coherence and empirical plausibility of results in relation to their domain of intended application.

2.2 Practical considerations

Despite this, there are some practical differences between modelling and simulation in the way that estimates are produced. Spatial microsimulation by IPF is more exacting as to the form of local area data used, requiring that the constraints be totals for categorical variables, commonly defined and for the same units of observation as in the survey cases. Modelling can make free use of scalar or ratio data, and local predictors need not be variables directly relating to the survey units of observation. Modelled estimates can make use of administrative sources like benefit claim rates or house sale prices, which pose considerable problems of consistency and definition for SMS. It follows from this that IPF, in its basic form, considers only household-level variance in income; inter-area differences are modelled only as inter-area differences in population composition. In

---

[2] The fact that income estimation is being attempted typically implies implies the absence of comprehensive income data against which the precision of any estimate might definitively be assessed

modelled estimates, inter-area differences in income can be conceived of as multi-level, including area area-level variables and efffects.

Modelled estimates however also demand that detailed spatial identifiers be available in the survey data, for the spatial units for which estimates are desired. In the UK context this is a significant hurdle, as access to survey microdata which identify the location of cases below regional level is, for sound reasons of respondent confidentiality, tightly controlled. Spatial microsimulation does not have this requirement, and thus can be carried out, as in this paper, from readily available public data sources; it does not require access to secured or commercial sources.

The other important practical difference between modelling and simulation lies in the form of the results produced. In modelled approaches, only a single statistic is modelled at one time, and correspondingly only a single statistic is estimated for each area. In the commonest models, this is mean income, and thus only area mean incomes are estimated. Logistic models might predict the probability of households being income poor, and thus estimate area proportions in income poverty; by means of quantile regression one can estimate medians and other points of local income distributions (Tzavidis et al. 2008). Each such specific income statistic which is wanted in the final results requires the specification of a separate regression model. By contrast, microsimulation weights, once generated in the re-weighting process, can be applied to the survey data to estimate, with relative facility, multiple income statistics for each area, and can even be used to project or test policy changes.

3 Setup, Sources and Methods

3.1 Linking variables

In the example here, SMS is used to estimate household income for the 33 boroughs (local authority areas) of London in the years 2001 and 2011. The survey source data is FRS and HBAI, and the local population data is taken primarily from the 2001 and 2011 UK Censuses of Population. This is supplemented with local administrative data on counts of dwellings by Council Tax Band; as discussed, other administrative data present problems of reconciling definitions and units of observation.

The specific variables in the local population data are chosen on the strength of being predictors of household income in the survey data. The selection of income predictors by regression methods has been treated systematically and in depth by both Anderson (Anderson 2007) and the ONS (Methodology Directorate and Centre for Regional and Local Statistics 2010). The present research drew on this work, and thus variable selection is not reported here in detail. The main differences in this paper is the use a set of variables common to both 2001 and 2011 datasets, and the possibility in the hierarchical fitting technique, described below, of using not only household-level but also adult-level Census and survey variables.

Table 1: Variables used to link the Family Resources Survey with the Census 2001 and 2011

Sim/ Num	Label	Measurement Level	Description
HHO/1	HHOLD.TYPE	Household	Household composition
HHO/2	EMPLOY.STAT.HRP	Household	Household representative's employment status
HHO/3	CTAX.BAND	Household	Council Tax Band of dwelling
HHO/4	TENURE	Household	Broad housing tenure
HHO/5	ACCOM.TYPE	Household	Type of dwelling
MLV/1	EMPSTAT.LIVARR	Adult	Employment status, by whether living in a couple
MLV/2	NSSEC.ACTIVE	Adult	NS-SEC Occupational class of economically active adults
MLV/3	ETHNICITY	Adult	Broad ethnic group
MLV/4	AGE.SEX	Adult	Sex and 10-year age group
SPI	(Taxable Income)	Adult Tax Payer	Approximated distribution of taxable incomes (SPI)

Table 1 summarises the linking variables used in the three variant simulations presented in this paper, the *Household-Only*, *Multi-Level* and *SPI Prior-Weights* simulations, which are shortly introduced. The variables are listed in reverse order in which the survey weights are adjusted to population totals, so that the variables listed first have, as it were, the highest priority. This simulated local area population totals for the last entered variable, household type, are thus always exactly correct. The source data, showing the derivation of the marginal constraints from census tables are available in the downloadable package accompanying this paper[3].

It is worth noting that there are here some minor discrepancies between the survey data and the populations to which they refer. The local population total should be the sampling frame for the survey in that area. However, for example FRS is a sample of private households only, whereas some Census tables include people living in communal establishments such as care homes, student halls and prisons. Similarly, Council Tax Band (CTB) records cover all dwellings, but unoccupied dwellings are outside the FRS's sampling frame. These discrepancies are tolerated as being relatively small[4] – indeed, CTB records are used in the FRS's own weighting scheme, despite the possible problems (Lound and Broad 2013, pp.9,17) – and are mitigated by ensuring that the highest-priority

---

[3] `https://github.com/a-fent/microsim-ipf`

[4] In 2011, 1.5% of London "usual residents" aged 16 or over lived in communal establishments, with the remainder in private households. Camden, where a large number of student halls of residence are located, has the highest proportion of residents (4%) in communal establishments.

variables cover wholly corresponding populations.

3.2 *The Household-Only simulation* HHO

The results in this paper derive from and compare three variants of the reweighting method. The first represents a standard IPF approach, and is referred to in the paper as the "Household-Only" (HHO) simulation, as it uses only household-level variables to link survey household cases with local population totals. In practice, this limits constraint variables to properties either of the household as a unit (its composition, number of children, tenure), of its dwelling (flat or house, tax band) or of its representative person (ethnicity, sex, age). The *HHO* simulation uses only the five variables listed in the topmost section of the linking variables in the reweighting procedure. It excludes some potentially useful household variables, notably number of earners per household, used in other work (Anderson 2007), since the necessary local population totals have not been produced for the 2011 Census.

3.3 *The Multi-level simulation* MLV

This restriction narrows the range of variables that might be used, since many Census totals are reported in more detail at the adult, rather than household, level. For example: in FRS, socio-economic class is reported only for economically active individuals, whereas the Census reports the NS-SEC of retired and unemployed people based on their last main job. Standard Census tables on class are thus not easily used as constraints as IPF, even though occupational class is, unsurprisingly, an important predictor of individual and household income. One must either, as Anderson does, make some judgement-based adjustment to the local totals of NS-SEC of household head (ibid., p.10 fn2), or forego the use of this predictor altogether.

Furthermore, a household-level-only simulation will produce estimated local totals for individuals that are inconsistent with actual known totals. Thus, for example, the age/sex structure of the population or individual employment rates are not guaranteed to be correctly reproduced, because there are systematic inter-area differences in adult circumstances not observable in household-level variables (for example, rates of employment for married women). This means that a source of potential inter-area difference is missed, reducing the dispersion of the estimates, and also makes the weights less plausible for use in further policy simulation.

The first refinement of IPF methods examined in this paper is thus a hierarchical or multi-level reweighting of the FRS to local population totals. In the *MLV* simulation, Census tables relating to adults *and* Census tables relating to households are used, as shown in table above. These are linked respectively to the FRS's `adult` and `household` datasets. This technique is sketched, but not extensively empirically explored, in a conference paper by Müller and Axhausen (Müller and Axhausen 2011). The procedure adopted

here follows that paper, whereby in each IPF iteration, the adult constraint totals are first applied. The arithmetic mean of these weights is then used as the starting weight to fit the household-level constraints. These weights are then in turn applied to all adult household members, and the adult constraints re-applied, and so on, for the desired number of iterations, finally fitting and producing a set of household weights.

3.4 Prior weights from tax income data: SPI

An open question in SMS is the specification of the seed or starting weights (Whitworth 2013, p.30). If each household is given a equal starting weight of, say, 1, as in standard IPF, it implies before constraining the survey data to the local population totals, we have no knowledge of the likely incidence of each household in each area. This may be a reasonable approach, but there may be grounds for using alternative starting weights. For example, when estimating incomes for London boroughs, we assume the starting weight of all London cases from the FRS is 1, but exclude FRS cases from all other regions, effectively assigning them a starting weight of 0.

The second refinement to IPF considered in this paper is additional adjustment to starting weights. A recurring difficulty in the research was in estimating top incomes, particularly in boroughs, such as Kensington and Chelsea and Westminster, where considerable numbers of extremely wealthy households live. This has consequences then for estimates of dispersion and inequality in the local income distributions. There is no apparent way to resolve this with only Census data, since pertitent categorical information, such as being in the highest occupational groups, identifies only a broad sweep of the better-off, rather than the very highest earners.

Therefore the third variant simulation uses the *Survey of Personal Incomes* (SPI) to set seed weights and supplement the categorical data used in the other two variants. The SPI data are derived from a large sample of tax return records, and thus provide annual income estimates that are reasonably accurate at borough level, although covering only taxable income of those who make tax returns. SPI data are used in official statistics to improve HBAI estimates estimation of very high incomes (Department for Work and Pensions 2014, p.242). Here the SPI is used to adjust the starting weight of each adult who is believed to have been assessed for income-tax (by having total non-benefits income greater than the personal tax allowance in the relevant year), such that the starting weight reflects the relative probability of an adult with such an income being drawn from the particular distribution of taxable incomes in that local area.

More formally, it assumes that the distributions of adult non-benefit incomes in each borough, and as sampled by the FRS in London, follow a log-normal distribution (Clementi and Gallegati 2005). The parameters of this distribution (its mean and standard deviation) in each borough and in London are calculated from the mean and median income figures in the SPI tables published by HMRC. The starting weight of each tax payer in each borough is the probability density of his or her income in the borough distribution,

relative to the probability density of the same income in the SPI distribution for London as a whole.

the seed weight of case i in borough b is:

$$seed.weight_{ib} = \frac{f(x_i; \mu_b, \sigma_b^2)}{f(x_i; \mu_L, \sigma_L^2)}$$

where $f(x; \mu, \sigma^2)$ is the probability density of the case's income, x, in a lognormal distribution with the parameters μ and σ^2. b and L denote the parameters of that distribution in a given borough and in London as a whole respectively. These weights are then scaled so that they sum to n, where n is the number of adult cases presumed to be taxpayers. The starting weights of adults who are assumed not to be income-tax payers, having a total non-benefits income of less than the Personal Allowance, is left unchanged at 1.

3.5 Convergence of the three simulations

IPF methods include a set of internal tests which check whether the reweighting procedure has correctly converged, that is, whether population totals using the reweighted dataset are consistent with the known population totals from which the weights are derived. The value of these tests here is primarily to confirm that the variant methods can produce formally correct solutions, and they are presented thus only briefly.

In the tests, the weights derived from the chosen number of IPF iterations (here, 20) are used to estimate totals of the link variables used in each area. These estimates are compared with the actual local population totals. A summary of this comparison by variable for each variants is given in Table 2 and Table 3. The metric used is the root-mean-squared-error (RMSE) - i.e. the square root of the mean squared differences between each cell in the original constraint table and the corresponding estimate got using the new weights. This measure "can be interpreted as an average difference in the true and estimated frequency count" (Simpson and Tranmer 2005, p.225).

Lower values are "better", and the final constraint applied – here, household type – necessarily has nil error. Beyond this, there are no clear prescriptions as to what values of RMSE and similar measures should be taken as acceptable (Anderson 2007, p.14). Here, a sense of the degree of error can be got by considering that London boroughs had in 2011 on average 100,000 households, and that the errors of the multi-level simulation are all less than 1,000 on the household variables.

We can also see that the *HHO* simulation has substantially greater error in its estimates of adult-level variables, which is as we would expect given that such variables are not controlled for in this variant. The other two simulations, which include adult-level constraints, have much lower absolute RMSE values for the adult variables, whilst increasing the error of the household variables, but by much lower absolute values. This confirms the technical feasibility of multi-level IPF with these data. It may be of interest to

Table 2: Root mean squared error of the constraint variables in the three simulations, 2001

fit.01	Constraint		Hhold Only	Multi-Level	SPI
	Areas	Classes	RMSE	RMSE	RMSE
Adult					
ETHNICITY	33	5	11,320	1,879	1,894
NSSEC.ACTIVE	33	9	5,049	738	734
AGE.SEX	33	16	1,959	308	311
EMPSTAT.LIVARR	33	11	3,852	554	566
Household					
ACCOM.TYPE	33	4	2	38	39
TENURE	33	4	2	135	149
CTAX.BAND	33	8	704	705	705
EMPLOY.STAT.74.HRP	33	5	11	29	32
HHOLD.TYPE	33	9	0	0	0

Table 3: Root mean squared error of the constraint variables in the three simulations, 2011

fit.11	Constraint		Hhold Only	Multi-Level	SPI
	Areas	Classes	RMSE	RMSE	RMSE
Adult					
ETHNICITY	33	5	14,686	2,156	2,150
NSSEC.ACTIVE	33	10	5,761	1,338	1,335
AGE.SEX	33	16	2,462	522	519
EMPSTAT.LIVARR	33	10	3,671	551	545
Household					
ACCOM.TYPE	33	4	0	15	16
TENURE	33	4	0	95	94
CTAX.BAND	33	8	627	627	627
EMPLOY.STAT.HRP	33	5	5	10	10
HHOLD.TYPE	33	7	0	0	0

note that in the single-level simulation, increasing the number of fitting iterations from 10 to 20 produced virtually no further reductions in error, whereas the error of all variables in the two multi-level simulations was further reduced by doing 20 iterations. Multi-level fitting, understandably, may require more repetitions of the computation steps.[5] The other point of interest is that the use of prior weights in the *SPI* simulation makes no appreciable difference to the quality of the solution.

4 Coherence and Plausibility of the Estimates

These tests of convergence indicate nothing about the value of the income estimates that can be derived from the SMS weights, although such tests are those which most widely appear in the literature. A recent paper goes so far as to argue that "attempts to validate the outcomes of [spatial microsimulation] are relatively weak, and much more strongly focused towards technical checks on the robustness and consistency of procedures than assumptions and outcomes" (Birkin and Clarke 2011, p.203). A simulation could converge perfectly and still not produce any meaningful estimates of income.

The very endeavour of SMS estimation, however, normally starts from of the absence of direct local observations of the variables of interest, against which the estimates might be tested. An alternative approach is firstly, to test the *coherency* of the aggregated London results against official published results, and then to check the plausibility of the modelled results using other, related, variables for which small area information is known(ibid., p.203).

4.1 The whole income distribution in London

Although HBAI does not provide borough-level income statistics, it does yield London-wide statistics from its standard weights. In the three simulations, London-wide weights for each case can be got by simply summing its weights in each of the 33 boroughs. With these, we can test the expectation that the London income distributions from the simulations are close to those in the official statistics, and that inconsistencies are explicable. Figure 1 presents the decile values of three household income measures: gross income, and income equivalised by the OECD scale before and after housing costs. It compares estimates using HBAI published weights, shown as boxes with confidence intervals, with estimates derived from the three simulations.

For all three simulations, for all three income measures, and at all deciles, the SMS estimates are coherent with the published HBAI estimates, lying within the 95% confidence interval of the HBAI estimates. The simulations of most interest, *MLV* and *SPI*

---

[5] The computational intensity of microsimulation is often remarked on. In a recent paper Tanton et al. report two other simulation techniques taking ½ and 2½ hours for 107 areas (Tanton, Williamson, and Harding 2014, p.85). Here, the more complex two-level fitting for 33 areas took about 12 seconds, parallelising the calculation of each area's weights by the `parallel` package included in recent versions of R.
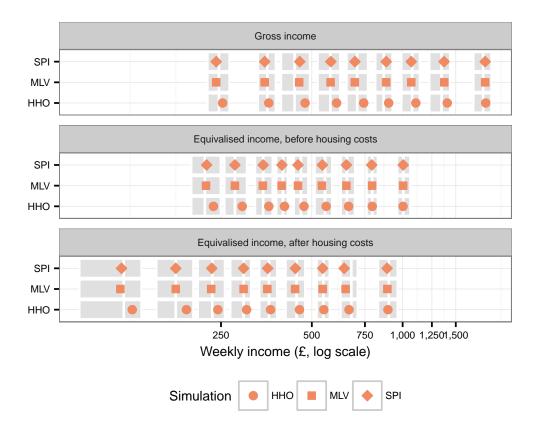
Figure 1: Estimates of decile values of 2011 household income in London (gross, and equivalised, before and after housing costs). Estimates from the published HBAI weights are shown as boxes with the 95% confidence interval. Points show the all-London estimates using the simulation weights. All estimates based on counts of benefit units.

lie especially close to the central official estimate. It is also useful to note that different measures of income (gross and equivalised, before and after housing costs) are coherently estimated from a single simulation. It may thus not be necessary in practice to conduct separate simulations to estimate specific income measures, and rather to "produce transferable models that represent multiple characteristics" (Birkin and Clarke 2011, p.197).

4.1.1 Regional accounts measures of household income

The only published UK statistics which offer a measure of household income below regional level are the previously mentioned tables of Gross Disposable Household Income (GDHI). These form part of ONS's *Regional Accounts*, and provide annual figures for regions, NUTS2 and NUTS3 areas, geographic divisions harmonised across Eurostat members. These tables give per-capita figures for "the amount of money that all of the individuals in the household sector have available for spending or saving after income
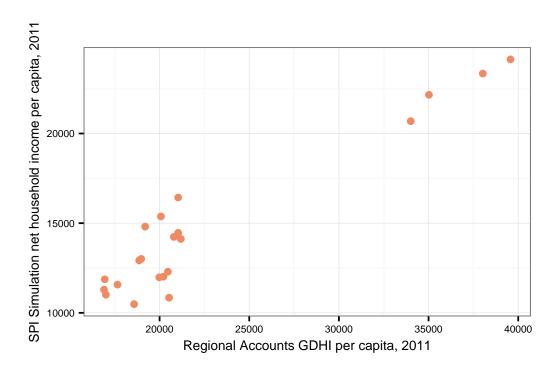
Figure 2: **Estimates of net household income per capita from the SPI simulation for London's NUTS3 (2015) areas, compared to ONS's published Gross Disposable Household Income per capita figures, from the Regional Accounts.**

distribution measures (for example, taxes, social contributions and benefits) have taken effect".[6]

Figure 2 presents household income in London's 21 current NUTS3 areas, which each comprise one or more local authorities. The chart compares the GDHI per-capita figures with a per-capita calculation based on the SPI microsimulation results (using the HBAI variable `HNTINCBU`, as the definition of net income closest to that in GDHI). Whilst a broad correlation is apparent, there are very considerable differences in absolute terms and in the relationship between the two measures in each area. This reflects a large gap in both method and definitions. GDHI is based on the downwards apportionment of National Accounts totals; its definition of "households" is correspondingly the "household sector", which includes non-profit institutions like trade unions and universities. The simulation results are bottom-up estimates based on survey microdata. Further reconcilation is not attempted here; it seems more important simply to note that GDHI, though not infrequently used in policy circles, is some considerable distance removed from SMS results based on HBAI, which are more clearly and closely tied to conventional notions of household income and welfare.

---

[6] `http://www.ons.gov.uk/ons/rel/regional-accounts/regional-household-income/`
`regional-gross-disposable-household-income--gdhi--2013/`
`stb-regional-gdhi-2013.html`

Table 4: Poverty rates (percentage of people in relative low income <60% of median) 2011, HBAI and simulations

Poverty rate (people)	HBAI	Household Only	Multi Level	SPI Adjusted
Before Housing Costs	15.8	14.6	16.2	16.1
After Housing Costs	28.6	27.5	29.3	29.2

Table 5: Proportion of adults aged 16+ living in a couple, 2011

Data Source	Published Statistics		Simulations		
	Census 2011	HBAI	Household Only	Multi Level	SPI Adjusted
In a Couple	0.48	0.57	0.53	0.48	0.48
Not in a Couple	0.52	0.43	0.47	0.52	0.52

4.2 Poverty rates and lower incomes

At the national level, one of the main uses of the HBAI data is the estimation of income poverty rates by reference to the household income distribution, as in the annual DWP publication, *Households Below Average Income*. Comparable poverty measurement at local level is an application of SMS which is of especial policy relevance and interest. Table 4 shows poverty rates for London in 2011/12 based on the published HBAI weights and on the simulation output weights. The poverty threshold used is a household equivalised income of less than 60% of the contemporary national median. The table shows that the ratios from the simulations are within ±1% of the HBAI value; this is well within the 95% confidence intervals around the HBAI estimate (compare Anderson 2007, p.14, Table 7). The SMS results appear then to be coherent estimators of poverty rates. Whilst the Household-Only simulation underestimates poverty compared to HBAI, the two multi-level simulations produce higher estimates.

Investigation shows the difference between HBAI and the MLV and SPI simulations to be principally attributable to the differences in weights and implied population totals single people and couples. The Multi-Level and SPI simulations explicitly control for living arrangements (whether an adult is living with a partner) by use of Census data. The grossing control factors used in the HBAI do not control for the marital status or living arrangements of adults, other than the number of lone parent households (Department for Work and Pensions 2014, p.250, Table A.24). HBAI cannot directly use Census data in its grossing scheme, since it must, as an annual survey, produce weights outside of Census years.

Single adults, especially those living alone, have higher non-response rates, which are not fully compensated for in HBAI's grossing factors. As table 5 shows, the published HBAI weights give estimates of the proportion of people living in couples which are

Table 6: Boroughs with highest and lowest poverty rates (after housing costs) under the three simulations, 2011

Poverty	Hhold Only		Multi-Level		SPI	
	sl.name	sl.val	ml.name	ml.val	sp.name	sp.val
Highest						
1	Newham	0.37	Newham	0.43	Newham	0.43
2	Hackney	0.36	Tower Hamlets	0.41	Tower Hamlets	0.41
3	Tower Hamlets	0.33	Brent	0.35	Brent	0.35
4	Haringey	0.32	Hackney	0.35	Hackney	0.35
5	Lewisham	0.31	Haringey	0.33	Haringey	0.33
Lowest						
29	Sutton	0.23	Bromley	0.22	City of London	0.22
30	Richmond u.T.	0.21	Sutton	0.22	Sutton	0.22
31	Havering	0.21	Bexley	0.21	Bexley	0.21
32	City of London	0.21	Richmond u.T.	0.20	Havering	0.20
33	Bexley	0.21	Havering	0.20	Richmond u.T.	0.20

much higher than the Census. Since single-adult households (whether pensioners or working age) have lower incomes, especially after housing costs, and, increasingly over the 2000s, higher poverty rates than households with two adults, the official HBAI estimates slightly underestimate the real prevalence of low income and poverty. The slightly higher estimates from the simulation may be regarded as equally or more accurate than the published HBAI results.

Also relevant is the dispersion of poverty rates by borough (analysis and interpretation of these results is taken up in detail in the companion paper). Table 6 shows the highest and lowest poverty rates (after housing costs) for London boroughs. As might be expected, the additional Census information used in the Multi-Level simulation and in the SPI simulation, such as individual ethnicity and living arrangements, substantially increases the estimates of poverty rates for the poorest areas, whilst leaving them little changed in the least-poor boroughs. The addition of tax income data in the *SPI* simulation changes the estimates of poverty rates hardly at all.

4.2.1 Deprivation indices and poverty proxies from administrative data

The introduction noted that deprivation indices have wide acceptance and use in policy-oriented analysis of the subregional and small-area distribution of poverty. The best

known are the *Indices of Multiple Deprivation*, a central component of which are estimates of the incidence of income poverty. These are derived fairly directly from administrative records of rates of receipt of means-tested benefits and tax credits.

Since such estimates are necessarily affected by eligibility for and take-up of such benefits, microsimulation estimates have been at various times mooted – and rejected – as a potentially more satisfactory method for estimating local income poverty. A comparison of the simulation results with poverty proxies based on administrative data is thus of interest. Figure 3 presents such a comparison. The dataset used is the *Unadjusted Means-Tested Benefits Rate* (UMBR), a publicly available dataset covering 2001 to 2013. UMBR provides a poverty proxy rate based on the receipt of major means-tested benefits against the mid-year estimated count of households resident in an area (Fenton 2015). The chart shows this UMBR figure for the 33 boroughs in London in 2001 and 2011, against the poverty rate estimated from the SPI simulations in those years.

In both years, there is the expected positive correlation between the UMBR proxy rate and the simulated poverty rate. It is, however, notable, that this correlation is much stronger and more consistent in 2001 (0.97) than in 2011 (0.84). This is consistent with existing regional analysis, which shows that benefit claim rates became decreasingly well correlated with regional income poverty rates over the 2000s (Fenton 2013a). Reasons for this include the increasing share of in-work poverty; administrative counts of receipt of out-of-work income-replacement benefits are not sensitive to inter-area differences in the extent low pay and high housing costs as a cause of poverty. A detailed discussion of this trend is found in the accompanying paper. Here, the pertinent findings are that the SMS estimates seem highly plausible in the light of administrative poverty proxies, but that the latter have become increasingly weak at identifying inter-area differences in poverty as conventionally conceived.

4.3 Upper incomes and income inequality

As outlined above, a shortcoming of basic IPF models is that they necessarily use only nominal or ordinal household variables to predict income from local population characteristics. In the basic simulations presented here, socio-economic classification was used as a predictor of income. However, this assumes that the incomes of persons in the same occupation and with the same characteristics are the same in all areas. There is no "area effect" on incomes; all inter-area differences in average incomes are a result of inter-area differences in composition.

There are several reasons that this might be an unsatisfactory assumption. For one, some contracts include standard additional payments to employees based on area living costs, such as "London Weighting". For another, there are variations in income within the broad socio-economic groups, such as "higher managers and professionals" or "semi-routine workers" , used in the basic model. We might anticipate that to some degree those with higher income use it to live in more expensive areas. In absolute money
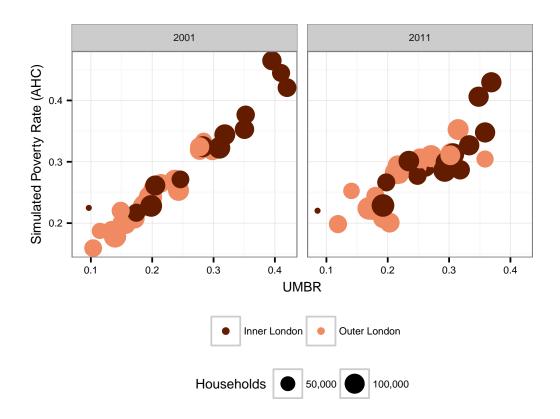
Figure 3: Comparison of poverty rates (income less than 60% of national median income, after housing costs) from the SP simulations with the Unadjusted Means-Tested Benefits Rate (UMBR) poverty proxy, by borough, 2001 and 2011.

terms, these differences in income will be greatest at the top of the income distribution, where a cadre of private professionals, executives and owners enjoy very high salaries and additional unearned income that distinguishes them from their more modestly paid peers in the same occupational group. This income elite in London is concentrated in boroughs like Westminster and Kensington and Chelsea (R. Atkinson 2015).

In an attempt to reflect some of this inter-area variation in income within ooccupations, the *SPI* simulation adjusted the relative weights given to tax-paying adults with like characteristics but different incomes by reference to income tax data for each borough. The refinement is anticipated to have greatest effect on the upper end of the resulting income distributions. To test this, estimates using the simulated weightings in each borough are compared to the *Annual Survey of Hours and Earnings* (ASHE), as an independent source with borough-level data. Using FRS and the SMS weights, borough medians of gross weekly earnings (the variable `INEARNS` in the `adults` table) are estimated, considering only full-time employees (`EMPSTATI == 1`). In figure 4 these estimates are compared to 2011 ASHE data on gross weekly pay at borough level, showing the range from the 40th to the 60th percentiles.
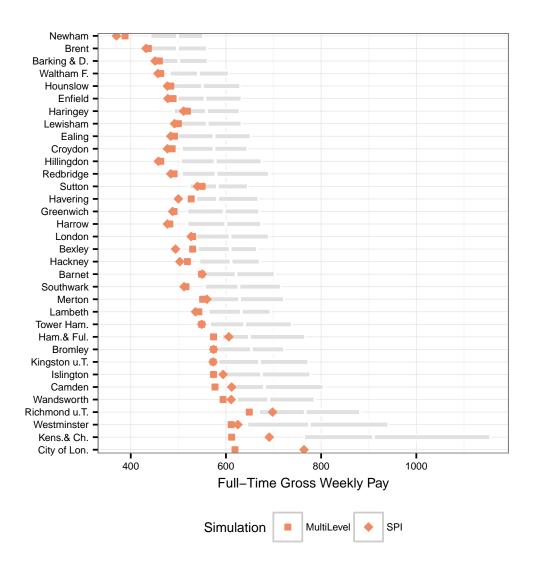
Figure 4: Point estimates of 2011 median gross weekly earnings from the multi-level (without tax data) and SPI (with tax data) simulations, compared to the 40th to 60th percentiles from published ASHE data, shown as pale grey bars. Data relate to full-time employees only.

In the majority of boroughs, the SMS estimates of the median lie just below the 40th percentile in ASHE. This is also true using the standard FRS weightings, reflecting differences in coverage between surveys.[7] However, there are large absolute disparities between the medians derived from the simulations which do not use the SPI data in boroughs with the highest incomes, most markedly in Kensington and Chelsea and the City of London. The additional use of SPI data moves the simulation median of earned income closer to the ASHE median in such boroughs – but at the same time, moves it further away in some boroughs with lower average earnings.

---

[7] ASHE is based on a much larger sample from employers' records, but excludes some casual workers and those paid less than the amount at which National Insurance becomes due. In comparison to FRS and other household surveys such as the Labour Force Survey, it covers lowest-paid workers less completely.

Table 7: Boroughs with highest and lowest Gini coefficients of income inequality (equivalised income after housing costs) under two simulations, 2011

Gini coefficient	Multi-Level		SPI	
	Borough	Gini	Borough	Gini
Highest				
1	Kensington and Chelsea	0.56	Kensington and Chelsea	0.61
2	City of London	0.56	Westminster	0.60
3	Westminster	0.55	City of London	0.59
4	Hammersmith and Fulham	0.53	Camden	0.57
5	Camden	0.53	Hammersmith and Fulham	0.56
Lowest				
29	Newham	0.37	Newham	0.36
30	Waltham Forest	0.37	Waltham Forest	0.35
31	Havering	0.36	Havering	0.35
32	Bexley	0.35	Bexley	0.34
33	Barking and Dagenham	0.34	Barking and Dagenham	0.32

Since the addition of SPI data increases the SMS estimates of upper income, we would expect it also to increase the estimated degree of income inequality in those areas. Table 7 provides a confirmation of this using the Gini coefficient. The greatest inter-borough income inequality is in the inner London boroughs with the highest top incomes, Kensington and Westminster, and the Gini coefficients increase with the adjustment using SPI data.

The results indicate that the adjusted starting weights produce somewhat more plausible results in a relatively small number of areas with very high incomes. However, aside from the additional data requirements and calculations, the technique has shortcomings, particularly the increased error of estimates in those areas, like Kensington, with a local income distribution that is extreme and distant relative to the all-London distribution sampled by FRS. This is because in the most atypical areas, a small number of FRS/HBAI cases are being used in the SMS to represent a large proportion of the resident households. In Kensington, in the SPI simulation, one single survey case represents over 4% of the resident households (by comparison, a typical household case in HBAI represents 0.05% of London households). It follows from this that estimates of inequality and top income from the SPI SMS are somewhat unreliable and identifying inter-year differences with confidence difficult.

This illustrates how SMS cannot compensate for deficiencies in the coverage of source

data. It is only filling in missing data, and where it does this starting from only few observations, as at the top of the income scale, it will perform poorly. Unfortunately, it cannot solve the numerous problems of poor statistical observation of income elites, which affects numerous other official statistics, such as ASHE.

4.4 Income-related household characteristics

So far we have considered the plausibility of income and poverty estimates using the SMS weights. A final test is whether and how well the simulation weights reproduce the distribution of other income-related household characteristics not explicitly controlled for, but for which comparable administrative data exist. This is of relevance also when assessing the potential usefulness of SMS weights for simulating the local effects of policy changes, where reasonably accurate counts of potentially affected populations are important.

The receipt of Housing Benefit (HB) is chosen here as a strongly income-related characteristic of policy interest, as a benefit that is available on a strictly means-tested basis to both private and social tenants in and out of work (Fenton 2013b), and also a policy subject to regular change and, since 2010, large cuts. Figure 5 shows the actual and simulated rates of receipt of HB by borough in 2011. The actual counts are DWP administrative data, and the simulated counts are the number of benefit units whose HB receipt is >£0 in HBAI (variable EHBENBU). They are expressed as rates per 1,000 benefit units.

As is normal, counts and rates derived by grossing up from survey-reported benefit receipt are much smaller than the corresponding totals in administrative data. It is the correlation that is of interest, and this (0.87; 0.93 excluding the small and atypical City of London) between the actual and simulated borough rates indicates that the simulation weights effectively, if not fully, estimate inter-area differences in HB receipt. This provides additional confirmation of the usefulness of the simulation in predicting low-income, and, potentially with further adaptation, in the simulation of policy changes.

5 Discussion

The paper started with the premise that spatial disparities have been most often discussed either in the form of *differences of aggregates* – Gross Value Added, employment rates and so forth – which obscure the primary and secondary distribution of income and its relation to income adequacy and well-being, or through the use of deprivation indices, whose methods and sources have tended to detach small-area research on poverty from the methods and terms used in national-level research. This paper extends existing UK work on spatial microsimulation by iterative proportional fitting, as a means for estimating local income distributions that can bridge the gap between these existing approaches. IPF is a fairly straightforward method, and simulations may be prepared, as in this paper,
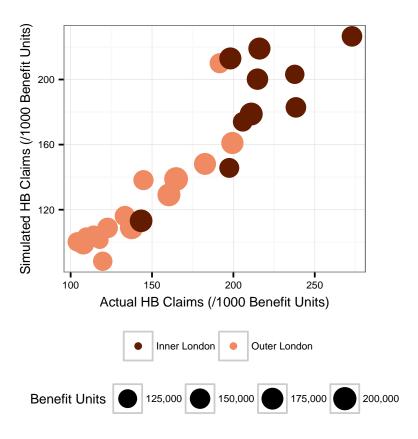
Figure 5: Simulated against actual housing benefit rates (claims per thousand benefit units), London boroughs, 2011. Figures from the SPI simulation shown against DWP's published figures for June 2011.

by use of no more than simple Census population totals and public-use survey datasets.

The analytic part of the paper showed that SMS produced local distributions of income that are coherent with the datasets from which they are derived, and which are plausible when compared to external sources on wages and benefits. One useful finding is that single simulations are capable of producing adequate representations of multiple domains, including different measures of income, its components, and income-related household characteristics like benefit receipt. This lends support to the view that for applied policy analysis, SMS can produce transferrable models that are also suitable for further use in testing policy scenarios.

5.1 The variant techniques

Two modest extensions to existing techniques were examined. The first of these, multi-level IPF, enabled simulation with local area and survey data measured at both the household and adult level, rather than only household data as in standard IPF. In this case, the primary effect of including adult-level data was refined estimation of incomes at the bottom of the end of the distribution. However, the main implications of multi-level IPF are

more general. Firstly, it expands the potential range of linking variables which might be used in future simulations - which is more relevant in non-Census years, when a much less rich range of local area totals are available. Secondly, it produces a set of simulation weights that reflect a local population much closer to "reality", thus enhancing the value of the local derived weights for use in policy testing.

The SPI simulation testing altering the starting weights of certain cases by reference to a scalar data source, the Survey of Personal Incomes. The primary effect shown was at the top end of the income distribution, especially in areas which have many highest-income households. That the simulation can be made to converge and produce plausible results when it includes SPI data suggests that these data could be incorporated into non-Census-based simulations. However, the technique means that estimates of top incomes become subject to very considerable uncertainty, which suggests caution about the use of standard inequality measures such as Gini with simulation weights.

5.2 Potential future developments

At the start of the last UK Coalition in 2010, the government returned to a well-established debate in official statistics, that measures of social progress ought to include more than just GDP. So far at least, this is in line with the argument made here that analysis of spatial disparities in economic wellbeing ought to attend, as far as possible, to intra-area distributions as well as inter-area differences. For the purposes of both policy and research, there would be considerable value in having a regularly updated and standard series of sub-regional household-income measures which include distributional information.

Nonetheless, there are obstacles to this. This paper has presented a set of consistent income simulations for two points in time (2001 and 2011), but these time points were dictated by the availability of rich census population data. Inter-censal simulations would need to draw on a wider range of sources, which each bring with them their own difficulties. One route is to make use of regular large-sample surveys, such as the Labour Force Survey, to provide the local population totals to which survey data are fitted. This however means fitting to survey estimates, which have their own (often quite large) uncertainty. Another route is projection of Census data forwards, as is done for the various official national population projections. Again, this introduces its own uncertainty, and when carried out from scratch is a labour-intensive process (Anderson 2007). Lastly more use may be made of administrative data, in the way that the Survey of Personal Income and Council Tax records have been used here and in the FRS's grossing controls. This then however adds considerable difficulties of incompatible definitions, under-reporting and differences in populations. Producing compatible survey variables and area counts is likely again to involve estimation and adjustment, even though we know, for example, that there are systematic differences in the survey reporting of benefit take-up between areas (Bramley, Lancaster, and Gordon 2000).

Whilst this paper has discussed simulated estimates at the level of large local authorities, there is considerable interest in and effort towards simulation and estimation at yet smaller spatial scales in Britain.[8] The obstacles described above to the production of local authority simulations outside of census years apply yet more forcefully to neighbourhood-scale geographic units. Supporting survey data from the Labour Force Survey, ASHE or SPI are of course not usable for such small units. A possible way forward would be to carry out two-stage simulation, first at local authority or county level with the benefit of survey estimates of key household and person characteristics, then at the smaller scale, using the upper-level weights as a starting point, and relying on a mixture of projected Census and administrative data at the smaller level.

This paper has largely left aside the potential (and historically important) use of IPF area weights in policy simulations. These would, for example, bring the kind of modelling of policy changes to taxes and benefits conducted by the Institute of Fiscal Studies at a national level to a local level, to show the spatially differentiated effects of fiscal policy (Ballas et al. 2005; Tanton and K. Edwards 2012). It is probably in this kind of applied research, as much as in the empirical analysis of results as taken up modestly in the accompanying paper, that spatial microsimulation can best make a case for a place in the applied social researcher's statistical toolbox.

6 References

Anderson, B. (2007). *Creating small area income estimates for England: spatial microsimulation modelling*. 2007-07. Colchester: CHIMERA, University of Essex.

– (2008). *Creating small area income deprivation estimates for Northern Ireland: spatial microsimulation modelling*. Colchester: Institute for Social and Technical Research (ISTR), University of Essex.

– (2012). "Estimating Small-Area Income Deprivation: An Iterative Proportional Fitting Approach". In: *Spatial Microsimulation: A Reference Guide for Users*. Ed. by R. Tanton and K. Edwards. Understanding Population Trends and Processes 6. Springer Netherlands, 49–67.

Atkinson, R. (2015). "Limited exposure: Social concealment, mobility and engagement with public space by the super-rich in London". In: *Environment and Planning A*, 0308518X15598323.

Atkinson, T. et al. (2015). *Monitoring the evolution of income poverty and real incomes over time*. CASE Papers 188. London: London School of Economics.

Ballas, D. et al. (2005). *Geography matters. Simulating the local impacts of national social policies*. York: Joseph Rowntree Foundation.

Betti, G. and A. Lemmi, eds. (2013). *Poverty and Social Exclusion: New Methods of Analysis*. Routledge. 313 pp.

---

[8] The ONS's small-area income estimates are for MSOAs; Anderson's various simulations, and also the GLA's recent income estimates go down to the smaller, and more widely used LSOA geography.

Birkin, M. and M. Clarke (2011). "Spatial Microsimulation Models: A Review and a Glimpse into the Future". In: *Population Dynamics and Projection Methods*. Ed. by J. Stillwell and M. Clarke. Understanding Population Trends and Processes Volume 4. Dordrecht, Heidelberg, London, New York: Springer.

Bramley, G. and S. Lancaster (1998). "Modelling local and small-area income distributions in Scotland". In: *Environment and Planning C* 16, 681–706.

Bramley, G. et al. (2000). "Benefit Take-up and the Geography of Poverty in Scotland". In: *Regional Studies* 34.6, 507–519.

Clementi, F. and M. Gallegati (2005). *Pareto's Law of Income Distribution: Evidence for Germany, the United Kingdom, and the United States*. Microeconomics 0505006. EconWPA.

Department for Communities and Local Government (2015). *Updating the English Indices of Deprivation: Government Response to Consultation*. London: Department for Communities and Local Government.

Department for Work and Pensions (2014). *Households Below Average Income: An analysis of the income distribution 1994/95 – 2012/13*. London: Department for Work and Pensions.

Dorling, D. et al. (2007). *Poverty, wealth and place in Britain, 1968 to 2005*. Bristol: The Policy Press. 125 pp.

Edwards, S. and S. Gault (2009). "Developing and maintaining microsimulation models at DWP". In: ESRC Microsimulation Seminar Series: Seminar 1. London.

Elbers, C. et al. (2003). "Micro–Level Estimation of Poverty and Inequality". In: *Econometrica* 71.1, 355–364.

Fenton, A. (2013a). "Should we be using social security benefits data as proxies for income poverty?" In: *Journal of Poverty and Social Justice* 21.3, 207–218.

– (2013b). *Small-area measures of income poverty*. CASE Papers 173. London: London School of Economics.

– (2015). *Unadjusted Means-Tested Benefits Rate (UMBR), 2001-2013*. Dataset. CASE, London School of Economics.

Gardiner, B. et al. (2013). "Spatially unbalanced growth in the British economy". In: *Journal of Economic Geography*, lbt003.

GLA Intelligence (2014). *GLA Small area household income estimates 2011 - 2012 - Methodology of experimental household income data model*. London: Greater London Authority.

Harrell Jr., F. E. (2015). *Hmisc: Harrell Miscellaneous*. In collab. with C. Dupont et al. Version 3.17-0.

Haslett, S. et al. (2010). "More for Less? Comparing small area estimation, spatial microsimulation, and mass imputation". In: *Section on Survey Research Methods*. Joint Statistical Meetings. American Statistical Association.

Lound, C. and P. Broad (2013). *Initial Review of the FRS Weighting Scheme*. Office for National Statistics.

Lumley, T. (2004). "Analysis of Complex Survey Samples". In: *Journal of Statistical Software* 09 (i08).

– (2014). *survey: analysis of complex survey samples*. Version 3.30-3.

Methodology Directorate and Centre for Regional and Local Statistics (2010). *Model-based estimates of households in poverty for Middle Layer Super Output Areas in England and Wales, 2007/08: Technical Report*. Office for National Statistics.

Müller, K. and K. W. Axhausen (2011). "Hierarchical IPF: Generating a synthetic population for Switzerland". In: *ERSA conference papers*.

New Zealand and Statistics New Zealand (2010). *More for less? using statistical modelling to combine existing data sources to produce sounder, more detailed, and less expensive official statistics*. Wellington: Statistics New Zealand.

Office for National Statistics (2005). *The 2011 Census: Initial view on content for England and Wales. Consultation document.* London: Office for National Statistics.

O'Hara, G. (2005). "A journey without maps: The regional policies of the 1964–70 British Labour Government". In: *Regional Studies* 39.9, 1183–1195.

Roe, D. and D. Rendle (2009). "Microsimulation at HM Treasury: methods and challenges". In: ESRC Microsimulation Seminar Series: Seminar 4. University of Sussex.

Simpson, L. and M. Tranmer (2005). "Combining Sample and Census Data in Small Area Estimates: Iterative Proportional Fitting with Standard Software". In: *The Professional Geographer* 57.2, 222–234.

Snorrason, H. et al. (2015). *Peer review report (United Kingdom) on compliance with the Code of Practice and the Coordination Role of the National Statistics Institute*.

Tanton, R. and K. Edwards (2012). *Spatial Microsimulation: A Reference Guide for Users: A Reference Guide for Users*. Springer. 272 pp.

Tanton, R. et al. (2014). "Comparing two methods of reweighting a survey file to small area data". In: *International Journal of Microsimulation* 7.1, 76–99.

Tzavidis, N. et al. (2008). "M-quantile models with application to poverty mapping". In: *Statistical Methods and Applications* 17.3, 393–411.

Whitworth, A. (2013). *Evaluations and improvements in small area estimation methodologies*. Methodological Review Paper. National Centre for Research Methods.

Wickham, H. and W. Chang (2015). *ggplot2: An Implementation of the Grammar of Graphics*. Version 1.0.1.

Wilcox, S. and G. Bramley (2010). *Evaluating requirements for market and affordable housing*. London: National Housing and Planning Advice Unit.