EUROMOD WORKING PAPER SERIES

EM 04/26

Geo-SWITCH: Simulating regional income distribution and policy impacts in Ireland

Karina Doorley, Iris Wohnsiedler, Richard O’Shea, Niall Farrell, Cathal O’Donoghue

June 2026

Geo-SWITCH: Simulating regional income distribution and policy impacts in Ireland[1]

Karina Doorley [a]

Iris Wohnsiedler [a]

Richard O'Shea [a]

Niall Farrell [a]

Cathal O'Donoghue [b]

[a] Economic and Social Research Institute (ESRI) & Trinity College Dublin

[b] University of Galway

Abstract

This paper introduces Geo-SWITCH, a spatial extension of Ireland's tax-benefit microsimulation model SWITCH that enables distributional impact analysis at the county level. We align nationally representative SILC microdata with county-level administrative benchmarks using small-area estimation techniques, comparing deterministic calibration and probabilistic (Conditional Monte Carlo) approaches. Internal validation favour the deterministic method, which we validate externally and adopt for a range of applications including quantification of sample variability. We map county-level disposable income, revealing a spatial gradient from higher incomes in Dublin and the commuter belt to lower incomes in the North-West, and incorporate bootstrap confidence intervals to quantify sampling uncertainty. We then simulate medical card eligibility, showing substantial regional variation in access to free healthcare that closely tracks income patterns. Finally, we assess Budget 2026 reforms, finding that permanent welfare increases are progressive across counties while the withdrawal of temporary cost-of-living supports disproportionately affects lower-income areas. Geo-SWITCH provides a flexible, validated tool for regional policy analysis in Ireland.

JEL codes: C31; R13; D31; H71; H75

Keywords: spatial microsimulation, Ireland, income distribution, income tax, welfare

Corresponding author:

Iris Wohnsiedler iris.wohnsiedler@esri.ie

---

[1] The analysis in this paper draws on the ESRI's microsimulation model, SWITCH. We are grateful to the Central Statistics Office (CSO) for providing access to the Survey of Income and Living Conditions (SILC) Research Microdata File, on which the SWITCH tax-benefit model is based. SWITCH is based on the EUROMOD platform. Originally maintained, developed and managed by the Institute for Social and Economic Research (ISER), since 2021 EUROMOD has been maintained, developed and managed by the Joint Research Centre (JRC) of the European Commission, in collaboration with EUROSTAT and national teams from the EU countries. We are indebted to the many people who have contributed to the development of EUROMOD. This work was carried out with funding from the ESRI's Tax, Welfare and Pensions Research Programme (supported by the Departments of Public Expenditure, Infrastructure, Public Service Reform and Digitalisation; Social Protection; Health; Finance and the Department of Children, Disability and Equality), which is gratefully acknowledged.

1 Introduction

Over the last two decades, regional inequalities have grown within many OECD countries. Inequality between regions can result in economic, social and political costs, ultimately damaging social cohesion and undermining trust in government [OECD, 2023]. Like many EU and OECD countries, Ireland has a nationally uniform tax and benefit system that is applied to household income, which is unevenly distributed by region. This implies that, by redistributing between rich and poor, the tax-benefit system also redistributes income between regions.

Quantifying the extent of regional income inequality and redistribution is important from a planning perspective. There is increasing demand for a more disaggregated breakdown of ex ante distributional impact analysis (DIA) - the effect of policy reform on the distribution of income - to inform policy decisions. However, the scarcity of survey data that is representative at a smaller than national level is a barrier to this type of evidence.

In many countries, including Ireland, information from the Census and other administrative sources provides aggregate statistics at the county level. However, these sources lack the rich household micro-level detail needed for DIA. We follow the growing spatial microsimulation literature [O'Donoghue et al., 2014, Tanton, 2018, Sologon et al., 2020] in using available administrative statistics at the regional level to augment survey data for Ireland in order to construct a dataset that is representative of the population at county level and suitable for DIA. We employ and compare two families of reweighting approaches -probabilistic and deterministic- to align data from the Survey of Income and Living Conditions (SILC) for Ireland to county-level totals on demographic, employment, education, welfare receipt and income aggregates from the Census and other administrative sources. Our aim is to preserve the detail of the underlying microdata while reflecting the aggregate population composition at a spatially disaggregated level. We work at the county level for two reasons. First, county totals are readily available across the constraint domains we require, enabling both reweighting and validation. Second, counties are a salient and actionable unit for policy design, resource allocation, and accountability in Ireland. We link the geo-referenced survey data to SWITCH, the ESRI's microsimulation model for Ireland, in order to allow for ex ante DIA at the county level. Recognising the increased call for uncertainty quantification in microsimulation [Goedemé et al., 2013], we use a bootstrapping method to incorporate variance measurement into the model.

Our contribution to the Small Area Estimation and Spatial Microsimulation literature is threefold. First, we apply and compare two broad classes of spatial microsimulation methods —deterministic and probabilistic reweighting— within a single policy context using the Irish microsimulation model SWITCH. We evaluate their performance side by side. Second, we provide confidence intervals for the regional distribution of income to illustrate the statistical uncertainty surrounding our point estimates resulting from spatial reweighting. Third, we apply our preferred methodology to study the regional distribution of income in Ireland, simulating the spatial impact of social policy changes. We investigate two salient applications which assess the regional distribution of free medical care and the regional impact of policy changes announced as part of Ireland's Budget 2026. While it is common to apply small-area estimation and spatial microsimulation methods to estimate the distribution of a social or economic outcome at the small area level [e.g. Elbers et al., 2003, Ballas et al., 2005, Campbell and Ballas, 2016, Morrissey et al., 2008], much less work exists to integrate spatial microsimulation approaches with tax-benefit modelling infrastructure. This paper builds on work of O'Donoghue et al. [2013], O'Donoghue [2021] and Kilgarriff et al. [2016], and others, who link spatial microsimulation techniques with a tax-benefit system in Ireland. We develop this platform to integrate spatial microsimulation with the SWITCH model, the ESRI's tax and benefit model based on the EUROMOD platform.[1]

The paper is structured as follows: Section 2 provides an overview of the relevant literature and introduces the two main families of methods we compare, deterministic and probabilistic. Section 3 introduces the nationally representative survey data as well as the administrative aggregate statistics at county level. Section 4 then describes the methodologies, their application and internal and external validation. Section 5 provides results from the application of our preferred spatial microsimulation approach. Section 6 concludes.

2 Small Area Estimation and Spatial Microsimulation Literature

Small Area Estimation (SAE) and Spatial Microsimulation (SM) are distinct yet related methodological approaches to estimate economic or social statistics at sub-national levels. This is achieved by linking microdata, that is not representative at a smaller spatial disaggregation level, and aggregated statistics that, while available for various spatial scales, lack the household and individual detail required for many research questions. Despite its central role in many applied fields, the SAE and SM literature remains fragmented, reflecting the demands of multiple disciplines—including geography, statistics, and economics—without full integration across them. Moreover, the type of microsimulation modelling required for DIA introduces distinctive requirements and challenges that are not always considered in methodological discussions of SAE and SM.

2.1 Small Area Estimation

Generally speaking, SAE procedures first estimate the conditional distribution of an outcome variable (such as income) using survey microdata. The spatial distribution of this outcome variable is then predicted according to the census distribution of predictors. Under specific distributional conditions, outlined below, these models give insight into the conditional expectation of the outcome accompanied by an estimate of precision. This method was first proposed by Elbers et al. [2003] and has since been employed in estimating the spatial distribution of poverty in over 60 countries [Elbers and van der Weide, 2014]. Bedi et al. [2007] give an overview of applications in developing countries, while Simler [2016] apply these

---

[1] The design and validation of EUROMOD and SWITCH are extensively described in Sutherland and Figari [2013] and Keane et al. [2023] respectively.

methods to the European context. Further applications of note include the charting of spatial poverty incidence in Ecuador [Hentschel et al., 2000] and Vietnam [Minot, 2000].

SAE methods provide insight into an outcome of interest, such as the poverty rate, at the small-area level. However, this is a single statistic for the outcome: it does not produce a profile of spatially-representative individuals or households which can be linked to a model of taxes and transfers. This is required to facilitate distributional impact analysis. Spatial Microsimulation methods can provide this insight. We turn our attention to this branch of the literature next.

2.2 Spatial Microsimulation

Spatial Microsimulation involves the application of policy changes to the spatially-representative profile of individuals or households. In contrast to the poverty mapping methods described before, this enables the flexible analysis of changes to the policy environment and the application to a wide range of outcome variables. This is achieved by sampling from national-level microdata, which can be linked to microsimulation models, according to known distributions at the small-area level. Several literature reviews provide an overview of spatial microsimulation methods and applications. Rahman et al. [2010] focus on deterministic reweighting methods, providing an overview of the generalised linear regression approach (GREGWT) and discussing methodological issues. O'Donoghue et al. [2014] survey spatially focused microsimulation models and the methodological choices across domains such as demography, welfare and transport planning. Tanton [2014] outlines methods and a framework for static and dynamic spatial microsimulation and, in Tanton [2018], highlights future directions, including the provision of uncertainty measures (e.g. confidence intervals) alongside point estimates.

2.2.1 Deterministic Methods

The deterministic methods we discuss and use are commonly referred to as calibration estimators. They share a core idea: adjusting initial survey weights as little as possible while ensuring the calibrated weights reproduce a set of benchmark totals derived from external information. Formally, they solve a constrained optimisation problem that minimises a chosen measure of distance between calibrated and pre-existing weights, subject to meeting the benchmark totals (often up to a specified numerical tolerance). Methods differ in the particular distance function and algorithm used, and whether additional features—such as upper and lower bounds to prevent extreme weights—are imposed. They are deterministic because, once the inputs (benchmark totals, starting weights, and any tuning choices such as bounds and tolerances) are fixed, the resulting calibrated weights—and hence the estimates—are uniquely determined by the procedure, with no stochastic component.

In their seminal work, Deville and Särndal [1992] and Deville et al. [1993] develop a unified calibration framework by presenting a family of distance measures and cases, where several widely used deterministic reweighting procedures are special cases within this framework. They differ mainly in (i) the dimensionality of constraint variables and (ii) the chosen distance function. The three deterministic approaches most relevant in our small-area estimation setting (IPF, GREGWT, reweighting) are therefore best seen as closely related variants within the same calibration logic and are described in the following.

Iterative Proportional Fitting (IPF) is a multivariate raking algorithm that fits within the calibration framework. It iteratively calibrates to benchmark totals in two-way or multi-way Tables[2]. IPF has been used in a range of relevant applications. Wiki et al. [2023] apply IPF to reweight subjective survey data on well-being so that it is representative at the small-area level in New Zealand, addressing the absence of local well-being indicators. Tomintz et al.'s SimSALUD model implements deterministic methods (among them IPF) tailored to health-related research on small-areas in Austria. This online resource is one of few accessible tools, illustrating that there is a gap in accessible, easy-to-use methods for spatial microsimulation. Panori et al. [2017] develop a spatial microsimulation model for municipalities in Athens, enabling an assessment of the spatial impacts of austerity measures. They report an excellent internal fit and a reasonable external validation. Finally, Lovelace et al. [2015] review the origins and use of IPF and offer practical recommendations for evaluating static spatial microsimulation models. They emphasise internal validation techniques, noting that external validation is often highly context- and data-dependent.

GREGWT was developed by the Australian Bureau of Statistics [Bell, 2000] and is grounded in the calibration framework of Deville and Särndal [1992][3]. In practice, it adapts the Deville and Särndal calibration solution with a tailored algorithmic formulation, using a bounded Chi-squared distance function[4] [Rahman et al., 2010, p.11]. GREGWT has been widely applied in Australian static microsimulation. NATSEM's (National Centre for Social and Economic Modelling) spatial microsimulation model, spatialMSM, is built around the GREGWT reweighting algorithm [Vidyattama and Tanton, 2010]. Vu and Tanton [2010] use this GREGWT-based microsimulation approach to analyse the distributional impacts of an Australian household stimulus package at both national and regional (Statistical Local Area) levels, highlighting substantial regional variation in who benefited. Using the same methodological approach, Miranti et al. [2016] examine regional inequality in disposable household income across two Australian states (NSW and Victoria).

Reweighting. Finally, reweighting survey microdata to national benchmarks is often carried out using the logit calibration distance proposed by Deville et al. [1993]. The key advantage of this approach is that it allows explicit lower and upper bounds on the calibrated weights, limiting extreme weights. Conceptually, it is within the same calibration framework as IPF and GREGWT.

2.2.2 Probabilistic Methods

Probabilistic approaches start from a different premise: rather than enforcing exact agreement with auxiliary benchmarks through calibration constraints, they aim to approximate the benchmarks as closely

---

[2] "The solution [...] in this case can be obtained by carrying out (until convergence) the classical raking ratio algorithm of Deming and Stephan (1940), sometimes called iterative proportional fitting." [Deville and Särndal, 1992, p.381]

[3] Bell describe their approach as "[...] a particular case of a more general class of calibration estimates introduced by Deville and Särndal" [Bell, 2000, p. 3].

[4] Implementing Singh and Mohl [1996] which is based on Deville and Särndal [1992].

as possible via stochastic, “intelligent” search or resampling [Rahman et al., 2010, p. 14]. In practice, households are repeatedly selected and/or resampled so that the resulting synthetic population matches the constraint variables in expectation. These approaches are flexible: constraints can be prioritised, cross-tabulated, and extended to distributional targets. Moreover, they do not fail to converge since a solution can always be produced - at the expense of failing to meet internal benchmarks. The trade-offs are higher computational cost, sensitivity to tuning and algorithmic choices, and the need to report the stability of results across runs through standard errors and confidence intervals.

Simulated Annealing. This is a stochastic optimisation technique whose name and intuition stem from thermodynamics: as a material cools, it gradually settles into a low-energy (stable) state. In spatial microsimulation, the “energy” is typically defined as a goodness-of-fit measure capturing how far a candidate synthetic population is from the auxiliary benchmarks. The algorithm starts from a random selection of households for each small area and then repeatedly proposes small random changes—such as swapping one household in the synthetic pool for another—accepting improvements but also, especially early on, occasionally accepting worse solutions. This controlled acceptance of worse moves helps the search escape local minima and explore the solution space more broadly. As the ”temperature” decreases, the algorithm becomes increasingly conservative and converges towards a configuration that provides a close match to the benchmarks while preserving the joint structure of the microdata. This selection of households with replacement can make the method computationally intensive. Examples of implementations are Germany’s dynamic microsimulation model *MikroSim*, where simulated annealing is used to adjust the synthetic population for each municipality [Münnich et al., 2021, p.246]. Comparative evidence suggests it can outperform alternatives like hill climbing [Whitworth, 2022] and it has been applied to disaggregate UK health data with detailed validation [Wu et al., 2022]. However, simulated annealing is computationally demanding—sometimes to the point of being infeasible under a large set of constraints—and its performance is sensitive to modelling choices and parameter settings.

Quota Sampling and CMC. In the Irish context, [Farrell et al., 2010] introduced quota sampling, a hybrid methodology developed for the SMILE (Simulation Model of the Irish Local Economy) spatial microsimulation model. Building on simulated annealing, it reduces the search space by avoiding evaluation of all possible household combinations: as the simulation proceeds, households that satisfy quota requirements are “locked in” and are no longer replaced. The approach also mitigates non-convergence and difficulty in meeting tight constraints by allowing constraints to be relaxed when necessary (for a detailed overview see O’Donoghue et al. [2013]). This algorithm has been used for a wide range of applications, including analysis of the impact of natural endowments on farm incomes [Haydarov et al., 2024], health applications examining how individual characteristics shape hospital utilisation [Morrissey et al., 2013], and work linking employment income data with spatial commuting patterns [Vega et al., 2017].

The quota sampling methodology has been extended by Farrell [2024] to provide multiple population snapshots via random sampling; this approach—termed Conditional Monte Carlo (CMC) sampling—selectively samples households to quantify uncertainty in the resulting synthetic populations. It has also been applied to measure socioeconomic vulnerability to flooding at the small-area level in Ireland [Farrell and Ceolotto, 2024].

2.2.3 Method Comparison

There are relatively few studies that directly compare deterministic and probabilistic methods within a single setting. Although these approaches start from very different premises, they can be evaluated side by side using the same validation metrics. Direct comparisons remain uncommon but those available are discussed here.

Rahman et al. [2010] offer a conceptual overview contrasting deterministic reweighting (GREGWT) with probabilistic approaches such as simulated annealing. They note that standard probabilistic implementations typically do not provide built-in measures of random sampling uncertainty (e.g., standard errors), a limitation that more recent work addresses through Conditional Monte Carlo (CMC) approaches that generate multiple synthetic population realisations [Farrell, 2024]. For deterministic reweighting, Rahman et al. highlight the risk of non-convergence: while tolerance parameters can be adjusted, there is often less flexibility to relax or prioritise constraints.

Tanton et al. [2014] provide a detailed empirical comparison of GREGWT and simulated annealing using Australian survey and census data. They encounter non-convergence for GREGWT in some cases and permit early termination of the probabilistic routine for feasibility. In terms of internal validation, simulated annealing generally outperforms GREGWT, while external validation is satisfactory for both methods with no major differences. They also compare weight distributions, noting that simulated annealing can yield more zero weights, implying reliance on fewer donor households. This may make downstream microsimulation outcomes more sensitive to a limited number of observations.

Both deterministic and probabilistic approaches ultimately rely on a spatial homogeneity assumption [Tarozzi and Deaton, 2009]. In our setting this means treating households observed in a nationally representative survey as informative for a given county once we have selected or weighted them using covariates that are believed to drive the outcome of interest. This limitation guides our selection of what outcomes can be credibly analysed.

3 Data

SILC Survey Data The core input microdata we georeference are drawn from the Central Statistics Office's Survey on Income and Living Conditions (SILC). Although SILC is a household survey, it is linked to administrative records on current income and social transfers which constitutes an important difference to the harmonised EU-SILC dataset. Moreover, it contains a more extensive set of individual and household characteristic variables. These data underpin the Irish tax–benefit microsimulation model SWITCH [Keane et al., 2023] and are collected and calibrated to be *nationally* representative. While SILC records each household's county of residence, sample sizes and non-response adjustments are not designed to support reliable inference at the sub-national level; hence our use of small-area estimation to georeference the microdata. The geographic identifiers nonetheless enable us to draw region-specific subsamples[5] for the sampling process (regional subsetting). We use the 2022 SILC wave to align with the 2022 Census, providing consistent county-level administrative totals as described below. We aggregate the SILC data to the household level and work at this unit when constructing synthetic county populations; households are kept intact and not split across units.

Constraint variables Constraints are chosen to ensure that the reweighted microdata is representative of the social-welfare recipient population and the income-tax base. This selection mirrors the selection of constraints used in calibration of the national SWITCH model (see Keane et al. [2023] for details of the national re-weighting methodology), with minor adjustments reflecting data availability at the county level and to allow the incorporation of distributional targets for the probabilistic method, such as income medians. The choice of final constraints is informed by three criteria: the availability and quality of county-level measures, the degree to which variables can be mapped cleanly onto SILC concepts, and their relevance for analysing tax–benefit incidence (see section A.4 in the Appendix for regression and random forest results to test the association with disposable household income.). Table A.1 provides an overview of sources and variable descriptions, as well as the corresponding SILC survey variables used as constraint variables.[6] The final set of constraint variables compromises demographic characteristics (age-sex bands), educational attainment, gross household income, welfare receipt for major benefit categories and labour market status (number employed, unemployed).

For household income, mapping administrative county data to SILC is less straightforward than for demographics because income definitions and sources differ. We require county-level control totals (or medians) under a comprehensive definition that includes income from welfare receipt. The CSO's 2022 Geographical Profiles of Income in Ireland provide such measures at county, administrative county, and ED (Electoral district) level. These cover all household income matched to administrative sources, including social welfare and education grants[7]. Table A.1 lists the SWITCH income concepts aggregated to compute each survey household's median gross income (row 3, column 4). This includes income from employment and self-employment including pensions and benefits.

For welfare receipt, we impose a set of constraints to ensure the calibrated sample represents welfare recipients, in line with the variables used for the national SWITCH calibration. Administrative counts from the Department of Social Protection's Annual Statistical Reports are mapped directly to the SILC administrative welfare indicators. The welfare payments used are Contributory pension, Illness benefit (contributory), One-Parent Family payment, Working Family payment; Carer's allowance; Jobseeker's assistance and benefit; State pension (Non-Contributory); Widow's pension (Contributory); Invalidity pension and Disability allowance.

---

[5] Figure A.1 in the Appendix shows the 26 counties used as the unit of observation (in black), corresponding NUTS2 regions as well as the additional 5 administrative county subdivisions.

[6] Variables employed for external validation purposes—rather than as constraints—are introduced in section 4.3.

[7] For more information see the Central Statistics Office release Geographical Profiles of Income in Ireland 2022.

4 Georeferencing SWITCH

In this section, we introduce and compare two small-area methods (one each from the deterministic and probabilistic families), providing a description of the setup[8], implementation and validation results. Both methods ultimately lead to a set of household weights for each county which ensure the data aggregates to the constraint totals outlined in the previous section 3. Both the deterministic and probabilistic methods use the non-georeferenced SILC household microdata and reweight or resample households for each county to create an independent synthetic population.

4.1 Deterministic Method: Calibration

Our main deterministic approach uses calibration by adjusting the initial survey weights to satisfy county benchmark totals while minimising deviations from the original weights.

We implement the deterministic reweighting using the *calibrate* function from the `survey` package in `R` [Lumley, 2020], applying the logit distance function as proposed by Deville and Särndal [1992], Deville et al. [1993]. We also implement an IPF-based version, but we present results for the logit distance function as our baseline because this is consistent with the calibration used in the national SWITCH model and alternative distance functions yield highly similar results. The logit distance function imposes lower and upper bounds on the calibrated weights, which helps avoid extreme values.

As starting survey design weights, we use the SILC *euroweight* variable. These weights have already been calibrated by the CSO from the original design weights to match national population totals by age–sex, NUTS3 region, and household composition. We run the calibration separately for each county and, unlike the probabilistic approach, when it converges, it matches all county-level benchmark constraints exactly (up to a tolerance that is set to 10^{-7}). We discuss convergence issues related to adding more constraint variables in the internal validation section.

4.2 Probabilistic Method: Conditional Monte Carlo Sampling

The probabilistic approach we use to generate small-area estimates is the so-called *Conditional Monte Carlo* method (CMC), which employs rejection sampling to select households conditional on satisfying constraint variables relevant to the target outcomes. The method originates in the quota-sampling framework developed by O'Donoghue et al. [2013], and has since been extended by Farrell [2024] for estimating poverty rates and in Farrell and Ceolotto [2024] for producing small-area income estimates in the context of flood-risk assessments in Ireland. We adopt CMC for three main reasons. First, it is highly flexible, allowing the sampling procedure to be tailored to specific applications, and it can be adapted to yield a feasible solution where deterministic methods may fail to converge under restrictive constraint sets. Second, it facilitates the incorporation of additional distributional targets—such as medians or other moments of constraint variables—beyond simple totals. Third, the approach has been developed and

---

[8] A more formal description can be found in Appendix section A.2.

applied in the Irish context, providing relevant benchmarks for our implementation. In what follows, we describe adaptations to the methodological setup in Farrell and Ceolotto [2024] to our specific set of constraints and the requirements of our application.

The CMC method uses a rejection-sampling algorithm to build a county population that matches benchmark totals. We first expand the SILC household pool by replicating households in proportion to their *euroweight* survey weights, then repeatedly draw a candidate household at random from the weighted pool and accept it only if it does not exceed the constraint totals, otherwise it is rejected and a new household is drawn. In addition to matching benchmark totals, we also impose a distributional target—the median gross household income. Households that meet the other constraints are additionally rejected if accepting them would push the county median above the threshold (plus a tolerance, set to €1,000). Household are drawn with replacement, meaning that a household can be assigned more than once to a given county's pool of accepted households. Crucially however, once an instance of a household is accepted it cannot be later rejected or swapped for a different household. This differs from Simulated Annealing, which allows later state changes. This makes the procedure computationally efficient and feasible at scale, but it may fail to satisfy the constraint targets exactly, as misallocations cannot be corrected ex post.

So far, this is in line with the setup of CMC as delineated in Farrell [2024] and Farrell and Ceolotto [2024]. However, a number of modifications are made to improve performance in our setting. First, to reduce computational burden, we draw a working subsample of candidate households and refresh this subsample from the full pool once it is exhausted; in our application, 30,000 candidates balanced speed and stability of results. Second, at the initial stages of the algorithm, we restrict candidate draws to households with children to aid convergence towards the child-related constraint totals - this is in accordance with the prioritisation of households with children in Quota Sampling of the SMILE model [Farrell et al., 2012, pp.111].

To prevent the algorithm from cycling on a small set of duplicate households, we track what households are chosen on each refresh of the small sample pool. If the number of selected households is below a certain threshold and the same households are repeated then we relax the constraint set in a deterministic schedule, allowing households to re-enter the set of feasible households.[9]

Relaxing the constraints and incorporating adaptations for cases in which the algorithm stalls or the original problem is infeasible leads to an enforced solution in cases where the deterministic method may fail to converge.

---

[9] If the algorithm stalls, we progressively relax constraints: after 5 stalls we drop education-related/child tallies (keeping employment, age–sex, income, welfare); after 10 stalls we also drop welfare (keeping employment, age–sex, income); after 15 stalls we drop income tallies (keeping employment, age–sex); after 20 stalls we turn off median tracking; and after 25 stalls we retain only employment constraints. We use 32 iterations -which took approximately twelve hours - based on managing computational intensity, while differences between iterations are very small.

4.3 Validation

Before integrating the reweighted data into the microsimulation model SWITCH, we assess the performance of each method in terms of internal validation, i.e. the fit of the synthetic populations against the variables used as constraints. Furthermore, we study the plausibility of the resulting weights, with a focus on changes in the weight distribution relevant for microsimulation. Lastly, we externally validate our chosen method against county-level totals for variables not included in the constraint set.

4.3.1 Internal Validation

In the internal validation, we assess how closely each method reproduces the constraint totals (and medians) used in the reweighting process.

For the deterministic calibration approach, the optimisation problem is defined so that county-level constraints are matched exactly, subject to the chosen distance function and bounds on the weights. As a result, the algorithm either fails to converge or, if it converges, attains an exact fit to all imposed constraints. In our application with the sample of households at hand, calibration converges for all counties, and the internal fit to the constraint set is perfect. To assess the robustness of this convergence for sensitivity to sampling, we also implemented the deterministic method using 1,000 bootstrap replicate weights. While calibration based on the original survey weights converges for all counties and achieves an exact match to the imposed constraints, convergence is not guaranteed for every replicate. On average, 98.7% of replicate calibrations converge successfully, with non-convergence arising in a small number of cases where the resampled weights render the constraint system infeasible (see Appendix Table A.4 for convergence rates by county).

For the CMC method, the termination of the algorithm is encouraged through heuristic adjustments, including refreshing the sampling pool and, where necessary, relaxing binding constraints. These features improve stability and feasibility but imply that the final synthetic populations do not, in general, match all constraint totals exactly. Instead, the CMC solution achieves an approximate internal fit, and we evaluate its performance by examining the distribution and magnitude of deviations from the target constraints across counties.

Table 1 summarises the constraint fit for the CMC method and the calibration method (the full set of constraints, including all age-sex brackets and income brackets, is reported in Appendix Table A.5). For each method, the Table reports the ratio of estimated to true, target values (ratio) and the Pearson correlation coefficient averaged across counties. As expected, given the problem formulation, the calibration method achieves an exact internal fit: by construction, the model–target ratio is equal to one for all constraints and counties, and the correlation with the target is correspondingly equal to one[10].

For the CMC method, only the employment constraints are fully satisfied in all counties (model values equal target values), as this was prioritised in the algorithm termination. Education constraints tend to be slightly under-achieved for upper-secondary education and over-achieved at lower levels as

---

[10] Up to a tolerance of precision up to 10^{-7}.

Table 1: Internal validation, Constraint fit summary by method

	prob. method (CMC)		det. method (Calibration)	
Constraint	ratio (model/target)	Pearson corr.	ratio (model/target)	Pearson corr.
Employment				
Employed	1.000	1.000	1.000	1.000
Unemployed	1.000	1.000	1.000	1.000
NILF	1.000	1.000	1.000	1.000
Age-sex brackets (averaged)				
Men (age brackets)	1.001	1.000	1.000	1.000
Women (age brackets)	0.999	1.000	1.000	1.000
Education				
In education	0.985	1.000	1.000	1.000
Primary	1.073	0.997	1.000	1.000
Lower-secondary	0.943	0.996	1.000	1.000
Upper-secondary	0.860	0.999	1.000	1.000
Third-level	1.094	0.999	1.000	1.000
Benefit receipt				
State pension (contr.)	0.994	0.999	1.000	1.000
State pension (non-contr.)	0.676	0.934	1.000	1.000
Widow's pension	1.003	0.993	1.000	1.000
Invalidity Pension	1.031	0.997	1.000	1.000
One-parent family payment	0.912	0.994	1.000	1.000
Working Families Payment	1.035	0.999	1.000	1.000
Jobseeker's Assistance	1.066	0.997	1.000	1.000
Jobseeker's Benefit	1.144	0.993	1.000	1.000
Carer's Allowance	1.059	0.999	1.000	1.000
Illness benefit	1.056	1.000	1.000	1.000
Disability Allowance	1.092	0.996	1.000	1.000
Income				
Median_income	0.959	0.987	1.012	0.996
Income brackets (averaged)	0.741	0.999	1.000	1.000

Notes. This Table reports the constraint fit for the probabilistic method (first two columns) and the deterministic method (last two columns). For the calibration approach, median income is not imposed as an internal constraint (while income brackets are constrained upon). However, for the CMC method, the median income is explicitly included as a calibration constraint. See Table A.5 for all absolute (average) values for the CMC method.

well as third-level education. The model also simulates too few households (ratios of simulated/target recipients of around 0.7), for non-contributory pension receipt. These results point to areas where further model refinement is needed. In contrast, the calibration method is able to match these constraints close to perfectly (columns 3 and 4).

While Tables 1 and A.5 provide an aggregate overview of the fit across counties we further examine the CMC results in more detail to assess whether misfit is systematic across all counties or concentrated in a subset of them. The scatterplot plotting correspondence between simulated and target values in Figure A.3 suggests that when discrepancies occur, they tend either to appear fairly consistently across most counties or to be limited to a small number of clear outliers (each point represents a county and its position relative to the diagonal indicates constraint fit). There is little evidence that a specific group of counties is inherently difficult to fit; instead, any misfit appears to reflect constraints that are challenging to satisfy simultaneously across all counties.

Overall, while the CMC method guarantees convergence and offers greater flexibility, the deterministic calibration approach achieves exact internal validation whenever convergence is attained. In our robustness analysis using 1,000 bootstrap replicates, non-convergence is found to be rare under the chosen constraint set. Given this and the advantage of perfect internal fit (up to numerical tolerance) when successful, we find that the deterministic method outperforms the probabilistic one in hitting the benchmarks perfectly.

4.3.2 Survey Weight Changes

Figure 1: Survey weight distribution by method, Gini coefficient

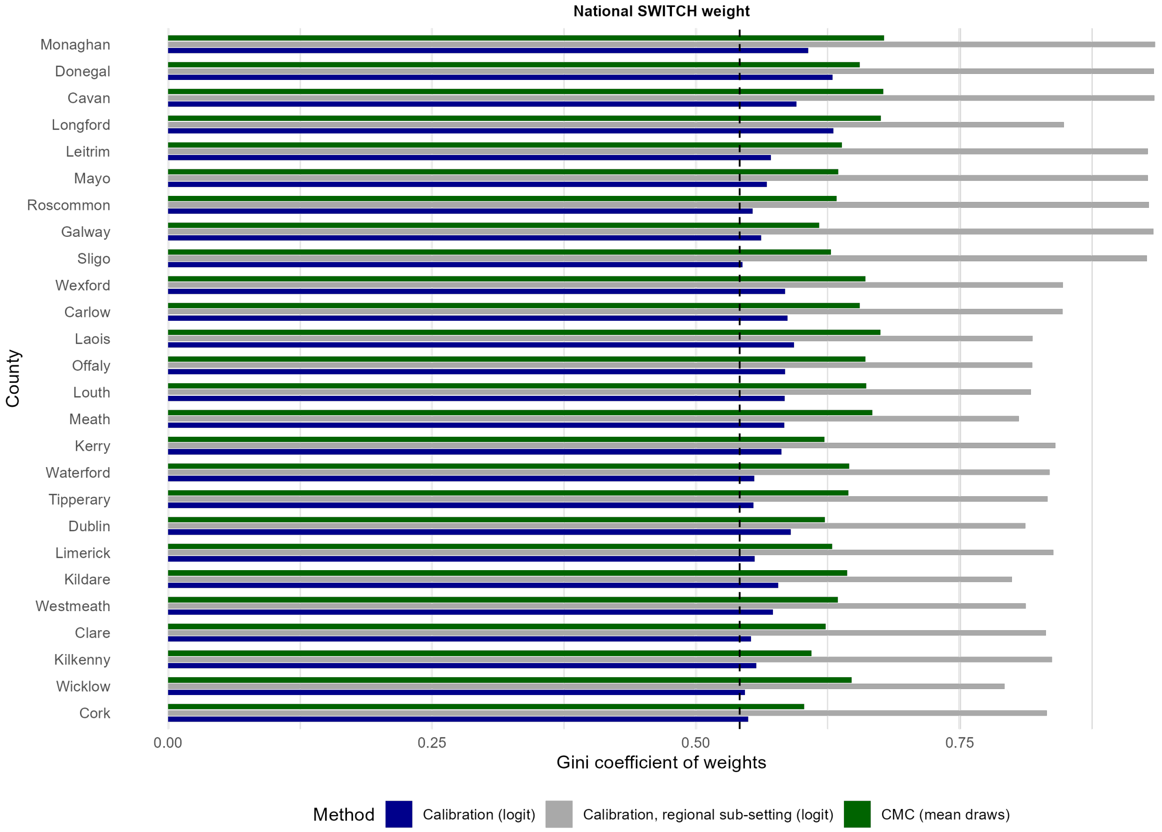

Notes. This Figure shows the concentration of survey weights by county and method, expressed as a Gini coefficient. A value of 0 indicates perfectly equal weights across households and a value close to 1 indicates that all of the weight is assigned to one household. The dashed horizontal line marks the Gini coefficient of the national SWITCH survey weights (0.542).

While CMC does not hit all constraints, another dimension to consider for validation is assessing how concentrated the resulting household weight distributions are. Survey weights are essential for making the surveyed households representative of the full population, but they can become problematic when a small number of households carry very large weights, making results more sensitive and unstable. This is a particular concern in spatial microsimulation, where the same set of survey households are used to construct subpopulations for all 26 counties and therefore tend to receive larger weights than for the national version. The literature highlights more extreme weights as a particular risk for calibration methods ([Deville et al., 1993, p.1020] and [Verma et al., 2017, p.7]), but they are not routinely reported.

For this reason, we include an assessment of weight distributions as part of the validation. This improves transparency for users and clarifies the types of applications for which the model can provide meaningful results.

For both methods, we compare the distribution of weights *before* and *after* adjustment: the "before" case is the national SWITCH weight; the "after" case is either the post-calibration weight (using the logit distance function) or, in the case of CMC, the number of times a household has been sampled. Figure 1 shows the concentration of weights by method (calibration and the average over CMC iterations), expressed as a Gini coefficient to measure inequality, where 0 corresponds to a perfectly equal distribution and 1 to the case where one household would account for the entire population in a county. There is no single "ideal" weight distribution, but more concentrated weights generally imply greater sensitivity of model outputs to individual households. A reference point is the weight distribution of national SWITCH model survey weights (dashed line). [11]

We also report results under regional sub-sampling (in grey, discussed in more detail in Appendix section A.5) to show how restricting the pool of eligible households to those within the same region affects the weight distribution. Under regional sub-sampling, weights are most highly concentrated, with Gini values approaching 1 (0.93 for Monaghan). This indicates that, in our setting, regional sub-sampling stretches the limited number of households too far.

The dark blue bars, which represent the calibration method (logit distance), are closest to the initial national distribution of weights. By contrast, the probabilistic method (CMC) exhibits a higher concentration of weights in all counties. In terms of the weight carried by the most heavily weighted households, under CMC the top 1% of households (ranked by weight) account for 17.1% of the total survey weight in a county, whereas under calibration they account for 15.4%. Similarly, the top 5% of households hold 41% of the total weight under CMC, compared to 36% under calibration.

These results indicate that the calibration method is more suitable for our use case, including when assessed in terms of the concentration of survey weights. However, relative to the national model, this still implies a higher concentration of weights, meaning that a smaller number of households carry more influence. This is an inherent downside of making the sample representative at a finer spatial level.

Overall, we conclude that in our county-level application - which is more aggregate than most small-area units used in the literature- the deterministic method converges with the chosen constraints and therefore outperforms the probabilistic method in terms of internal benchmark fit. We therefore proceed with the deterministic georeferenced households when assessing external fit.

---

[11] Which exhibit a slightly higher Gini than the initial survey weight, euroweight (0.542 compared to 0.526).

4.3.3 External Validation

External validation assesses how well the methods reproduce county-level quantities that are not directly constrained in the reweighting procedure. We focus on variables for which reliable county-level aggregates are available from administrative or Census sources, and consistent measures can be constructed from the SILC input data. This allows us to evaluate the capacity of our chosen method to generate credible estimates beyond the imposed constraint set. Table 2 summarises the external validation variables and the correspondence between the administrative benchmarks and the underlying SILC variables for deterministic reweighting.

Table 2: External validation, Overview of variables and sources

Category	Set	Variable(s)	Census 2022 Source	SILC [*SWITCH var.*] equivalent
Welfare	Primary	Share HHs with majority income from **welfare benefits**	CSO Geographical Profile of Admin. Income (2022b) (GPIIA11)	Transfers: [*ils_ben*] (public pensions, non- and means tested benefits), Market income: [*ils_origy*]
Welfare	Primary	Share HHs with majority income from **state pension**	CSO Geographical Profile of Admin. Income (2022b) (GPIIA11)	Pension: *ils_pen* (state pension, widows contr. pension, invalidity pension), Market income: *ils_origy*
Welfare	Primary & Secondary	**Child benefit** recipients, **Rent supplement** recipients	DSP Annual Statistics Report 2022	Administrative welfare receipt data from SWITCH
Income	Primary	HH **disposable income**	CSO Estimates of HH income (2022b)	Disposable income (SWITCH output): [*ils_dispy*]
Income	Secondary	Income from **self-employment** (total)	CSO Estimates of HH income (2022b)	Self-employed income: *yse*
Economic Status	Primary	Share **student, retired, unable to work due to permanent sickness or disability**	Population aged 15 and over usually resident and present in state (2022a, F7022)	*les* (6 student (15 and over), 4 pensioner, 8 sick or disabled)
Demogr.	Primary	**HH size**	CSO Number of families, persons or children in private households (2022a) (F3069)	Total number of persons by household
Housing	Secondary	Tenure type: share **renter/owner**	CSO Private households (2022a) (TAH05)	Owner: *amrtn=1,2* (owned outright, mortgage), Renter: *amrtn=3, 4, 6* (tenant, subtenant, rented at reduced rate, rented for free)
Housing	Secondary	Share HHs **living in house/flat**	CSO Private Households by Type of Accommodation (2022a) (SAP2022T6T1CTY)	HH010 (Dwelling type - 1,2 for detached/terraced house, 3,4,5 for apartment/flat/other)

Notes. All statistics are computed at county level. The CSO disposable income series (2022b) is not directly comparable to SWITCH/SILC disposable income; differences in coverage, definitions (e.g., equivalisation), and the treatment of taxes/transfers mean levels and growth rates diverge. Self-employment income (*yse*) in SILC/SWITCH is known to be volatile and may be under-reported; negative and zero values can occur and are retained, so comparisons with CSO aggregates should be interpreted cautiously. Children are defined as persons aged under 18. "Majority income" denotes households where >50% of gross income derives from the specified source.

For external validation, we classify targets into *primary* and *secondary* sets based on their relevance for tax–benefit simulations and the closeness of the mapping between administrative benchmarks and SILC/SWITCH concepts. Primary targets are those that are either core outcomes of interest or closely linked to them, and for which the benchmark-to-survey correspondence is relatively direct: the shares of households with majority income from benefits[12] and from state pensions; the shares retired and students

---

[12] Note that we treat the share of households for whom the majority of income comes from welfare benefits as an external variable because it is not used as a calibration constraint. While the model constrains on two categories (welfare benefits and public pensions; see Table A.1), this measure is constructed from the full set of welfare variables (e.g., non-contributory

(defined over the population aged 15+); the number of child benefit and rent supplement recipients; household size shares; and the key income aggregates (disposable income and market income). Secondary targets provide additional checks on fit along housing and compositional dimensions and include dwelling type (house/flat), tenure (owner/renter), self-employment income, and disability status.

Table 3: External validation, average summary statistics across counties

	Mean					Correlation coef.	
Variable set	target	simulated	Bias (units)	Rel. bias	RMSE (units)	Pearson	Spearman
Primary							
Share HHs majority benefits	0.303	0.298	-0.006	-1.703%	0.010	0.987	0.978
Share HHs state pension	0.151	0.142	-0.009	-5.546%	0.013	0.942	0.948
Share retired	0.167	0.168	0.001	0.997%	0.007	0.966	0.954
Share students	0.108	0.108	0.000	-0.309%	0.002	0.978	0.972
Child benefit recip. (abs.)	24,744	24,835	90	-0.328%	1,087	0.999	0.993
Rent supplement recip. (abs.)	348	331	-17	175.163%	584	0.930	0.663
Share in HH size: 1	0.089	0.072	-0.017	-18.798%	0.018	0.936	0.903
Share in HH size: 2	0.210	0.206	-0.004	-1.736%	0.006	0.957	0.951
Share in HH size: 3	0.192	0.214	0.022	11.467%	0.023	0.489	0.456
Share in HH size: 4	0.242	0.264	0.022	9.284%	0.024	0.831	0.732
Share in HH size: 5+	0.268	0.245	-0.023	-8.525%	0.026	0.840	0.858
HH disp. income (total)	5,194,834,615	4,138,795,346	-1,056,039,270	-20.395%	1,933,353,638	0.999	0.993
HH disp. income per capita	24,220	19,419	-4,801	-19.623%	5,030	0.639	0.504
HH gross median income	55,414	56,041	627	1.135%	838	0.996	0.999
HH gross income (total)	7,681,454,615	5,580,080,104	-2,101,374,511	-27.791%	4,016,052,869	0.999	0.996
Secondary							
Share HHs in a house	0.922	0.883	-0.040	-4.073%	0.055	0.825	0.636
Share HHs in flat	0.078	0.117	0.040	77.403%	0.055	0.825	0.636
Share income from self-empl.	0.114	0.089	-0.025	-20.852%	0.027	0.486	0.465
Share disabled	0.050	0.060	0.010	20.571%	0.010	0.975	0.954
Tenure type: Owner	0.708	0.677	-0.031	-3.987%	0.052	0.438	0.411
Tenure type: Renter	0.292	0.323	0.031	13.197%	0.052	0.438	0.411

Notes. This Table summarises external validation results at the county level. All statistics are averaged across counties, giving each county equal weight (i.e. results are not population-weighted). The Table therefore reflects the average model fit across counties rather than the fit for the average individual/household. Shares of retired persons, students, and disabled individuals are calculated relative to the population aged 15 and over. Child benefit and rent supplement are reported as absolute numbers of recipients. Household size shares refer to the proportion of the population living in households of size 1, 2, 3, 4, or 5+. Simulated county-level means are additionally averaged across bootstrap replicate weights, so reported estimates reflect the mean across bootstrap iterations and are thus robust to sampling variability.

Summary Results. Table 3 summarises results across counties by constraint (see also Table A.9 in the Appendix for an error summary heatmap by county). The first two columns, *Mean target* and *Mean simulated*, report the official ("true", typically Census-based) values alongside the simulated values from our small-area estimation. Bias (units) is the average simulated minus target in the variable's units (shares for most constraints; absolute counts for Child Benefit/Rent Supplement recipients or total Euros in case of the total household disposable income), indicating whether we over- or undershoot on average.[13] *Relative bias* reports the bias relative to the target value while the root mean squared error (RMSE) aggregates squared deviations and is a useful measure when many small errors coexist with a few large ones; in

---

state pension, child benefit, maternity benefit, education grants, rent supplement). It is then related to a household's market income (from employment and investment). Since this composite measure extends beyond the constrained categories, it is independent of the fitting procedure and therefore suitable for external validation.

[13] Note that results are reported unweighted across counties (i.e., not weighted by target or population size).

our setting it is typically larger than the mean absolute error. The final two columns report *Pearson* and *Spearman* correlation coefficients between simulated and target values, capturing linear and rank-based (monotonic) association, respectively.

For the primary constraints (those most closely linked to the model's core outputs), average bias is generally small and cross-county agreement is very high. The excellent fit for benefit and state pension shares suggests the model captures large benefit programmes well which is reassuring for our tax-benefit simulations. By contrast, the household-size distribution shows some systematic misallocation: larger household sizes (3 person or 4 person households) are systematically overstated. While the total population is well matched, its composition by household size shows some drift. This partly reflects the modelling choice to preserve household units, which restricts flexibility in reallocating individuals across household types.

We assess whether explicitly incorporating household size as an additional constraint improves overall model performance (Appendix Table A.8). Although the fit improves slightly for some external targets (e.g. share retired, disposable income) and mechanically achieves perfect fit for household size, it deteriorates for others (e.g. share students, state pension), and convergence rates decline substantially (87% on average compared to the 99% without the household-size constraint). Given the already satisfactory external fit under the baseline specification, and the considerably higher convergence rate, we retain the specification without household-size constraints.

For county-total disposable household income, the spatial ranking is matched almost perfectly, but the level exhibits a sizeable negative bias of about 20%, indicating a scaling issue rather than a model-fit one. This likely reflects definitional differences between administrative aggregates and SWITCH linked to coverage, components, and construction of disposable income.

For the *secondary* constraints, the fit is weaker also because these variables are less directly linked to the model's core outcomes and several (especially housing-related variables) may violate spatial homogeneity assumptions. County-specific and historical factors (urban/rural structure, housing markets) can limit how representative households are across space. It is reassuring that the shares of disabled people and students are captured well, although more volatile measures with high variance - such as income from self-employment - are not well aligned across counties. Overall, the external validation indicates a solid fit on key targets and spatial patterns.

Log Bias by County. Next, we assess fit by county to determine whether deviations from the targets are widespread or concentrated in a few outlier counties. One challenge with measuring bias by constraint and comparing this across all constraints is because the underlying variables have different units and distributions. To address this, we present the bias in the form of a log ratio bias measure, where bias $b_{c,j} = \log(\frac{\sum_{i \in N} w_{i,c} x_{i,j}}{T_{c,j}})$, with c indexing counties, j constraints, $x_{i,j}$ is the micro value for household i, $w_{i,c}$ are the calibrated weights, and $T_{c,j}$ is the target. A value of $b_{c,j} = 0$ indicates perfect fit; $b_{c,j} = 0.1$ means the simulated value is about $e^{0.1} - 1 \approx 10.5\%$ *above* the target, and $b_{c,j} = -0.1$ about 9.5% *below*. This metric treats over- and under-estimation symmetrically, is scale-invariant (counts, shares, income),

Figure 2: External validation, heatmap of bias, primary constraint set

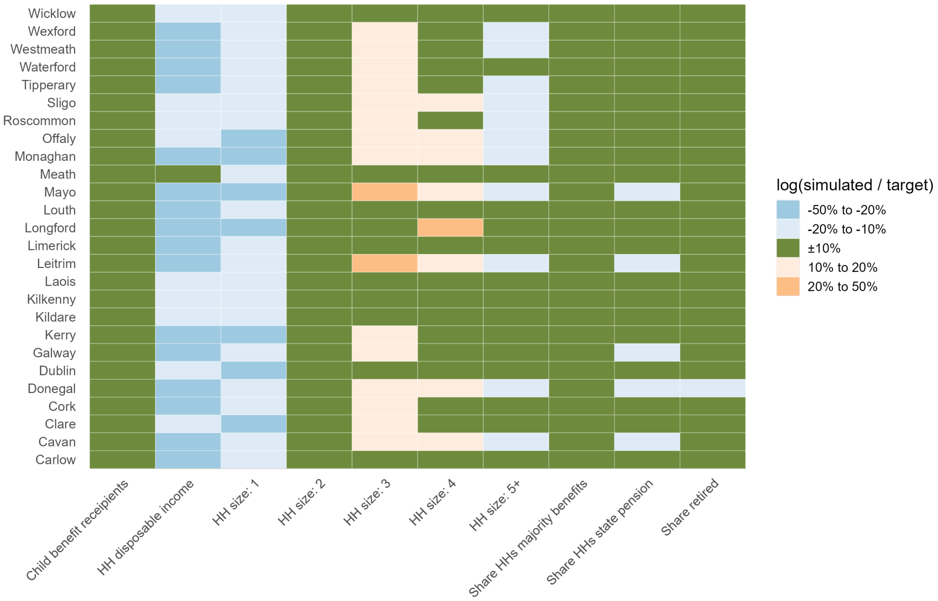

Notes. This heatmap shows correspondence between small-area estimates (calibration, logit distance function) and the primary set of constraints (as columns) by county (rows) in the form of a log ratio of the bias (actual/target).
Source. See Table 2.

is less sensitive to outliers, and can be re-expressed as a percent bias. For presentation, we plot $b_{c,j}$ in a county-by-constraint heatmap in Figure 2 for the primary set of external validation constraints.[14]

Unsurprisingly, the heatmaps echo the patterns in Table 3, but add the county dimension. For the household-size constraints, misfit is broad-based across counties with an underrepresentation of single-person households and a slight overrepresentation of 3 and 4 person households. For majority income from state pension, the overall correspondence is high, but there are clear pockets of mismatch—most notably Cavan, Donegal, Galway, Limerick, and Mayo—where the simulated shares diverges substantially from the targets (labelled in red). These discrepancies likely reflect more extreme and more variable demographic structures in these counties (see Appendix Table A.2 reporting Census statistics by county).

Map Comparison. Finally, to visualise the external validation spatially, we present a side-by-side map comparing the observed (target) and simulated mean values across counties, illustrating how closely the spatial patterns correspond. For brevity, Figure 3 shows results for the external variable share of households whose majority income comes from benefits — an external validation with strong relevance to our model outputs. As shown in Figure 3, the share of households whose majority income comes from benefits exhibits close correspondence across most counties. The main exceptions are Donegal,

---

[14] For the secondary set as well as the percentage deviation measure see Appendix section A.6.

Figure 3: Simulated vs. target values, share households with majority income from benefits

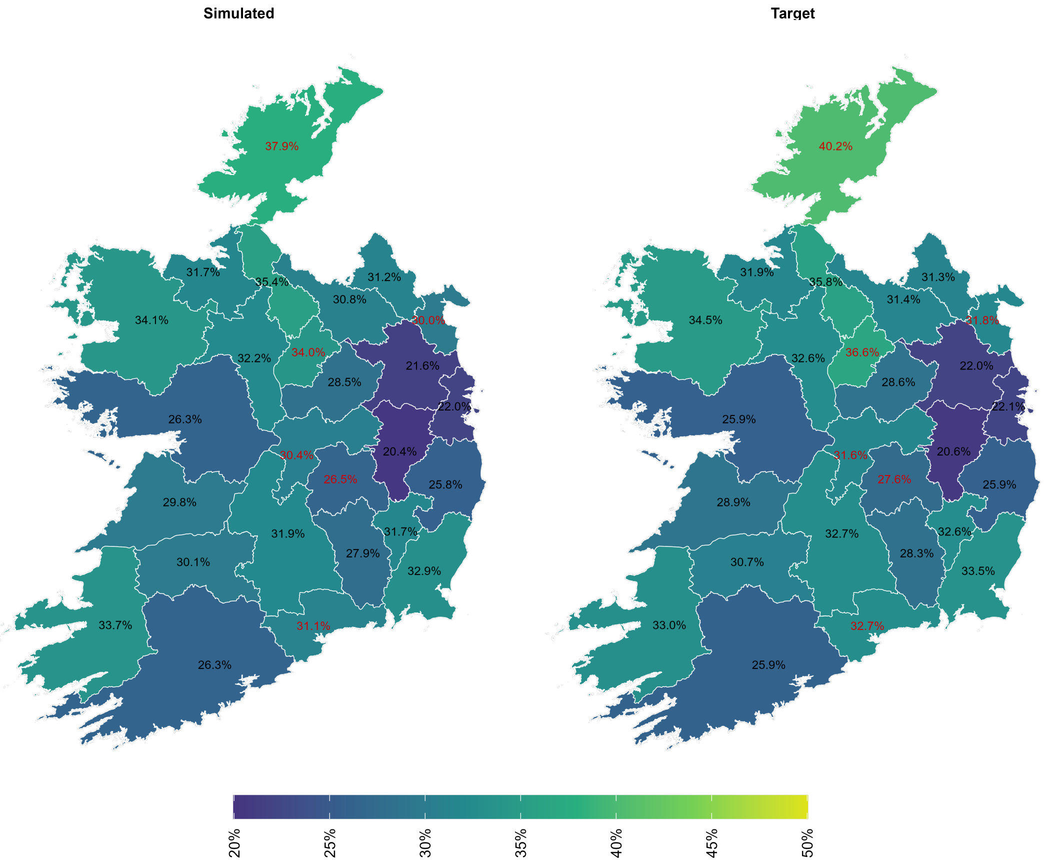

Notes. This Map shows simulated values for share of households with majority income from benefits. The left panel shows simulated values and the right panel shows target (census) values for each of the 26 counties. County labels appear in red where the absolute difference between simulated and target is more than 1%; all other labels are black.
Source. See Table 2, for county labels see Map A.1.

Louth, Longford, Offaly, and Laois, where the model slightly underestimates the share. Overall, this performance is reassuring, especially given that we did not explicitly target this metric, and it enables us to use the microsimulation model's flexibility in adjusting benefit eligibility and rates. Subject to the target outcomes, the external validation indicates that the model can produce meaningful county-level estimates for policy analysis.

4.4 Inference and Uncertainty Evaluation

Because our inputs are survey microdata, all statistics are subject to sampling variability, and our county-level reweighting introduces an additional source of uncertainty. To assess whether the resulting spatial point estimates are meaningful, and whether the method is reliable for policy analysis, we estimate variance measures using bootstrapping methods.

There is a general call for routine uncertainty quantification—not only in small-area estimation, but in microsimulation more broadly [Goedemé et al., 2013]. In addition, EU regulations mandate that poverty estimates be accompanied by measures of sampling variability [Tiziana and Ilaria, 2020]. For spatial microsimulation specifically, the "biggest difficulty" in evaluating the performance of small-area estimates has been described as the lack of such reliability measures (such as standard errors or confidence intervals) [Tanton, 2014, p. 80].

Given that the sampling design largely determines the variance, detailed information about the sampling process (stratification, clustering, weighting, and non-response adjustments) is needed. Hence, credible variance estimation requires access to design metadata or replicate weights provided by statistical agencies. The first-best solution would involve access to the initial design weights to create replicate weights and run them through the full sequence of non-response and calibration procedures [Wolter, 2007]. In the absence of this, Goedemé [2010], Goedemé et al. [2013] have shown that replicate weights after adjustment for non-response and calibration produce similar results. This is what we use to construct bootstrap replicate weights.[15]

Deterministic spatial microsimulation modelling techniques are largely drawn from methods of geocomputation to identify the combination of best fit subject to a set of constraints. Estimates of precision are not commonly identified using these methods. This has been acknowledged by Whitworth et al. [2017], who employ a multi-level regression model to prioritise constraint variables and use the between-area error term from this model to estimate a bootstrapped uncertainty estimate.

Bootstrapping creates a set of replicated estimates that we use in the microsimulation runs, and the variance of these runs captures both the variability in the parameter estimates and the variability due to model error [Klevmarken, 2022]. We illustrate this confidence-interval construction in our first applica-

---

[15] We use jackknife replicate weights supplied by the Central Statistics Office of Ireland, based on the final analysis weights (after nonresponse adjustment and calibration), to recover information on strata and primary sampling units. The Irish SILC uses a two-stage stratified design, with counties and income deciles as strata. Design weights are adjusted for nonresponse using logistic regression by wave, then calibrated to population totals by age–sex, tenure, household composition, and NUTS3 region.

tion in the next section. All mapped estimates are reported as the mean across 1,000 bootstrap replications, ensuring robustness to sampling variability in the reweighted microdata, and are accompanied by plots with confidence intervals.

5 Simulating regional level outcomes using Geo-SWITCH

In this section, we present the results of a series of simulations using the spatially adjusted data linked to SWITCH, hereafter Geo-SWITCH. Based on the comparison between methods carried out in the previous section, we use only the reweighting approach to adjust the SILC data, which is one of the commonly used deterministic methods. We first show how household disposable income is distributed by county. We then simulate entitlement to free medical care and illustrate how this entitlement is distributed across Ireland, with implications for demand and supply of healthcare. Finally, we show the regional effect of the full set of policy changes announced as part of Budget 2026.

5.1 *Geo-SWITCH* Simulations

5.1.1 Disposable Household Income

Figure 4 presents a choropleth map of county-level disposable income estimated by Geo-SWITCH for 2022. Each county is labeled with its median annual disposable income. The shading indicates whether the county median lies below (red) or above (blue) the national benchmark median income of €52,839.[16]

Figure 4 shows a spatial gradient in disposable income: Dublin and the surrounding commuter-belt counties are consistently above the national benchmark, while northern and western counties (especially Donegal, Leitrim, Mayo and Kerry) fall substantially below it. This pattern is consistent with the demographic composition of these counties (Table A.2) and the concentration of high-paying employment in and around Dublin.

Figure A.8 in the Appendix also shows how pre-tax and transfer (or market) income is distributed by county and how redistribution by the tax-benefit system varies at a regional level, with low-income counties as net beneficiaries and high-income households as net payers in the system.

We re-run the deterministic calibration for each one of the replicate weights[17], as discussed in section 4.4 and report associated confidence intervals. Figure 5 reports these at the county level for median disposable household income.[18] The confidence intervals lie relatively close to the point estimates, suggesting that sampling uncertainty is moderate and that the main county-level patterns in Geo-SWITCH are estimated with reasonable precision. Although confidence-interval widths vary somewhat across counties—implying, for example, more precise estimates for Meath than for County Dublin (Figure 5)—the

---

[16] The national benchmark is a weighted median using the nationally calibrated weight for SWITCH.

[17] Using the `R` package `survey` [Lumley, 2020] to approximate the sampling distribution via repeated, design-aware resampling (respecting stratification, clustering, and weights).

[18] Note that the results for the following applications are also available but not shown here.

Figure 4: Median household disposable income (2022)

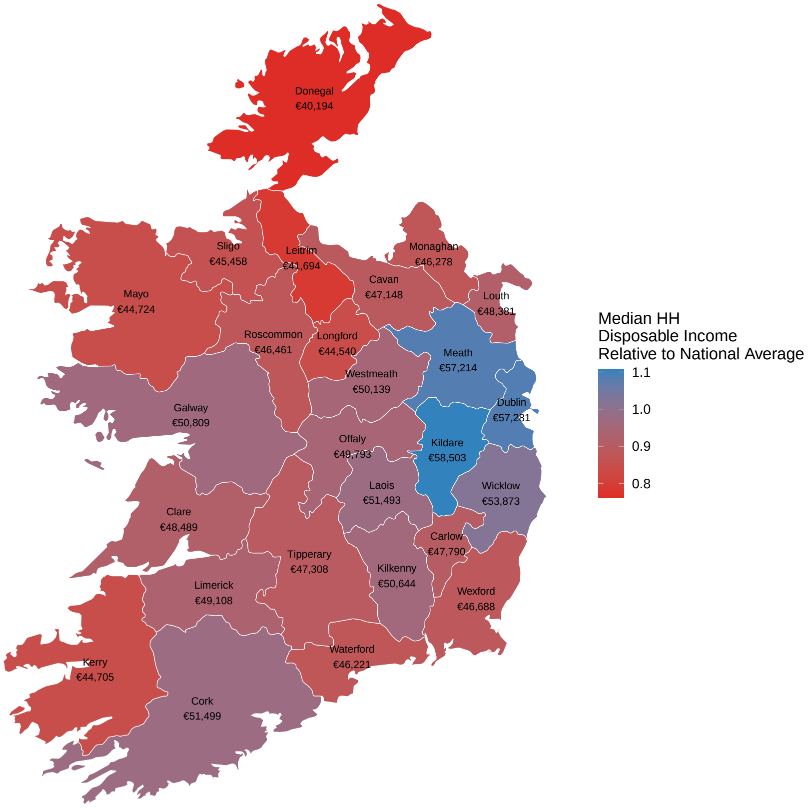

Notes. This map shows annual median household disposable income as of 2022 (mean as label). Colours represent each county's value relative to the national benchmark (normalised to 1.0) from the standard (non-georeferenced) microsimulation model. The legend is scaled to the observed range of county-to-national ratios, so 1.0 is not necessarily centred on the colour bar. This is the mean value across bootstrapped replications of the deterministic method to ensure robustness to sample variability.

differences are not large, and overall the county-level estimates are sufficiently precise to support policy-relevant interpretation. Note that while non-overlapping intervals can suggest meaningful differences between counties, this is not a formal significance test of pairwise differences. For a comparison of confidence intervals using regional sub-sampling, see Figure A.4 in the Appendix.

Figure 5: Median household disposable income (2022) with 95% CI

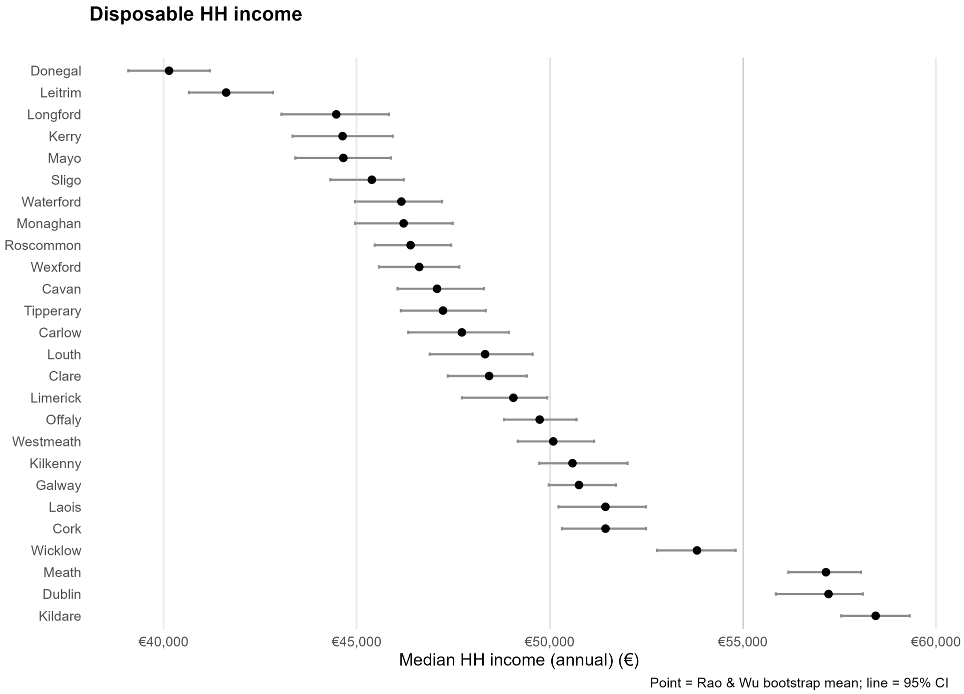

Notes. Points show mean county-level estimates based on the full sample using bootstrapping; horizontal lines indicate 95% confidence intervals [Rao and Wu, 1993]. Counties are ordered by increasing point estimate for median disposable household income (annual values in Euros, 2022).

5.1.2 Medical Card Eligibility

Medical cards are a major non-cash component of the Irish welfare state, conferring free (or heavily subsidised) access to a range of health services and reduced prescription charges. Eligibility is primarily means-tested, but the system contains important rule-based discontinuities and exceptions such as higher means thresholds for people aged over 70.

Assessing medical card eligibility using survey or other data requires microsimulation as it depends on disposable resources and eligibility rules interact with other elements of the tax-benefit system. Because entitlement affects behaviour and service use [Nolan and Layte, 2017, McDonnell et al., 2022], modelling entitlement at a granular level is important for policy makers to plan accordingly.

Figure 6 shows clear geographic variation in simulated medical card eligibility. The pattern closely mirrors the county distribution of disposable income for the same year: counties with higher average disposable income tend to have lower eligibility rates, consistent with the largely means-tested design of the scheme. In particular, Meath, Kildare and Dublin display the lowest shares of individuals eligible for a medical card. By contrast, eligibility is highest in more peripheral and generally lower-income counties—most notably Donegal and Longford—where the share is as high as 0.38. This spatial gradi-

Figure 6: Medical Card, share individuals eligible (2022)

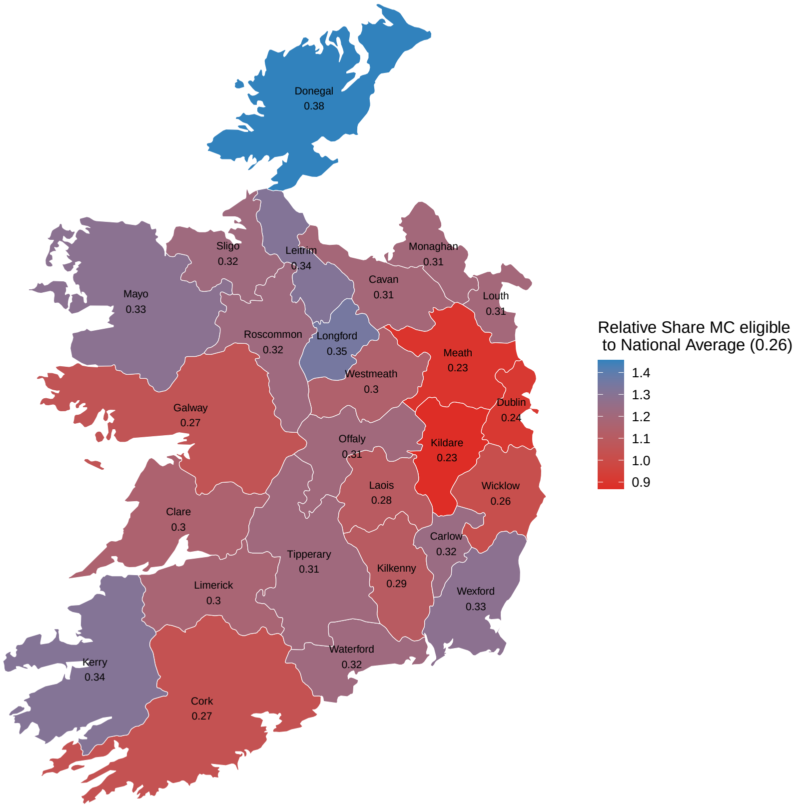

Notes. This Map shows the share of individuals eligible for a Medical Card (mean as label). Colours represent each county's value relative to the national benchmark (1.0) from the standard (non-spatialised) microsimulation model. The legend is scaled to the observed range of county-to-national ratios, so 1.0 is not necessarily centred on the colour bar.

ent highlights how differences in income and demographic composition[19] across counties translate into different exposure to, and reliance on, non-cash health supports. Estimated confidence intervals for this application by county are depicted in Figure A.9 in the Appendix.

5.1.3 Budget Impact Analysis 2026

We now turn to the regional impacts of Budget 2026 on disposable income. Budget measures are announced each October and typically apply to the following year. Microsimulation allows us to ex ante simulate the effect of the combined policy changes on household incomes.

The main policy changes announced in Budget 2026 - described in more detail in Bercholz and Simon [2025] - were a freeze to most tax thresholds and bands, corresponding to an effective tax increase

---

[19] For some key demographic and benefit/employment statistics at county level from Census 2022, see Table A.2

Figure 7: Budget 2026, Change in disposable household income, total effect

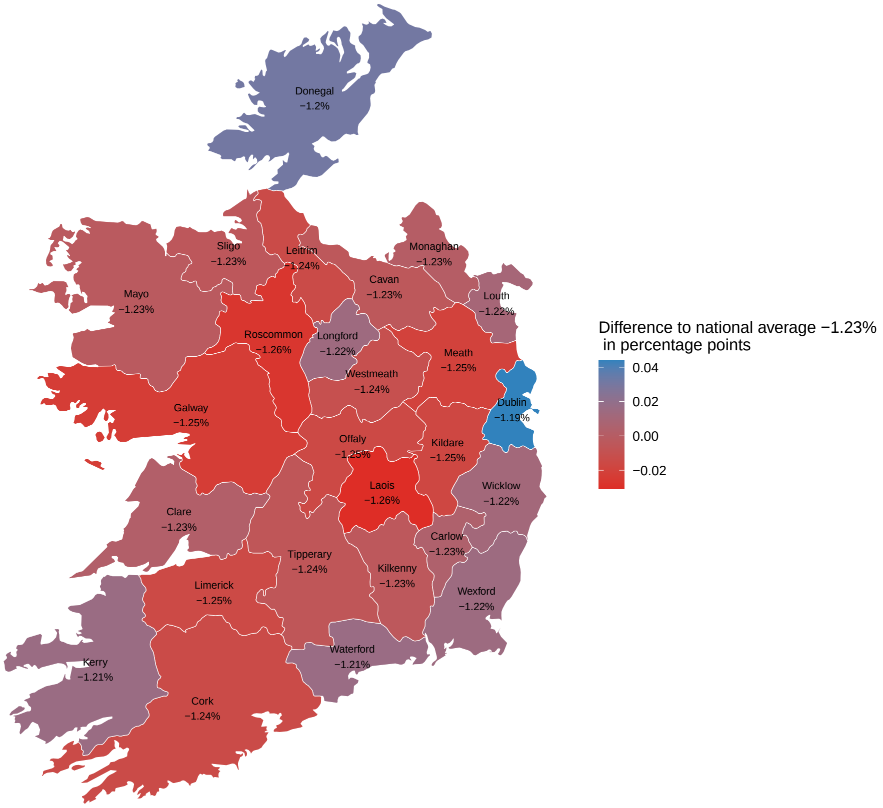

Notes. For each county, this Map reports the percentage change in mean disposable household income under the reform scenario relative to the price-uprated baseline scenario. Map shading shows the county's deviation from the national average (in percentage points: county change minus national change); blue indicates counties doing better than the national average and red indicates counties doing worse. Text labels show the percentage change between baseline and reform scenario. The national benchmark is computed using the original survey weight (*dwt*), while county estimates use the calibrated county weights derived from the deterministic method (averaged over bootstrapped replicate weights to account for sampling variability).

in the presence of rising earnings. On social welfare, there was a shift away from broad one-off-cost of-living payments, an increase to core weekly welfare rates of €10 and significantly increased child-related supports.

We compare Budget 2026 to an indexed baseline, a scenario in which the 2025 parameters of the tax-benefit system evolved in line with forecast price growth of 2.2% for 2026. Nationally, average disposable household income is estimated to be 1.23% lower under Budget 2026 than the indexed baseline[20]. Figure 7 shows each county's own percentage change in mean disposable income between the baseline and reform scenarios. The shading indicates each county's deviation from the national percentage change,

---

[20] -1.28% if we account for indirect tax changes, which we abstract from in this analysis.

Figure 8: Budget 2026, Change in disposable household income, temporary measures only

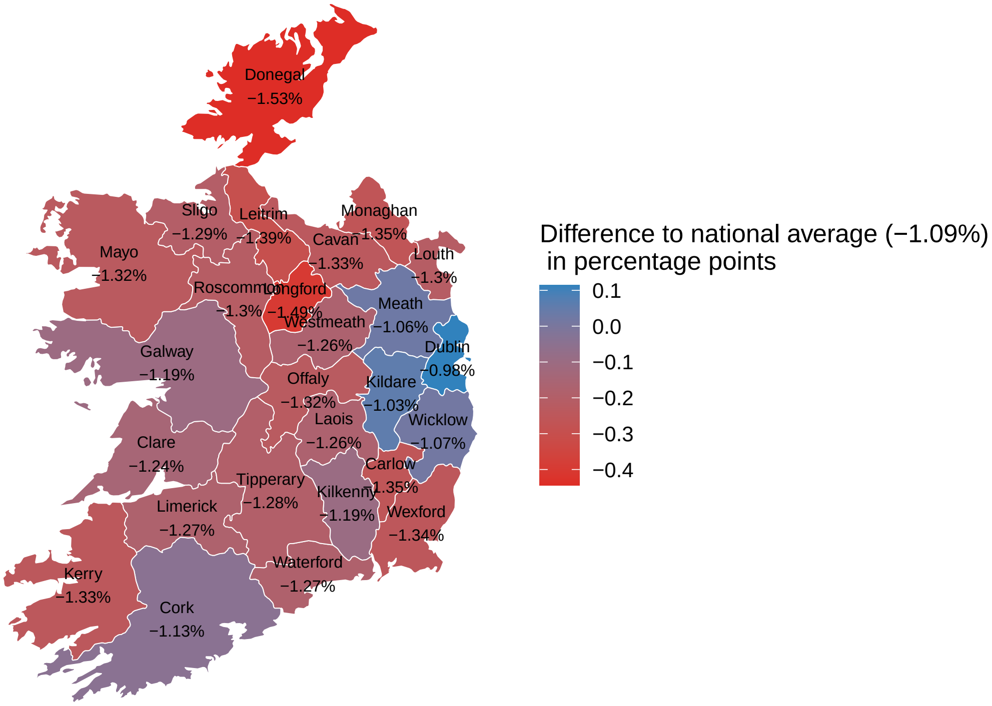

Notes. For each county, this Map reports the percentage change in mean disposable household income under the reform scenario (change in temporary measures only) relative to the price-uprated baseline scenario. Map shading shows the county's deviation from the national average (in percentage points: county change minus national change); blue indicates counties doing better than the national average and red indicates counties doing worse. Text labels show the percentage change between baseline and reform scenario. The national benchmark is computed using the original survey weight (*dwt*), while county estimates use the calibrated county weights derived from the deterministic method (averaged over bootstrapped replicate weights to account for sampling variability).

expressed in percentage points. County-level differences appear minimal but this masks heterogeneity in the effect of permanent changes to the tax and welfare system and the withdrawal of temporary measures. Figure A.10c in the Appendix shows the corresponding confidence interval plot.

The effect of the withdrawal of temporary cost-of-living supports and the introduction of permanent tax and welfare reforms are shown separately in Figures 8 and 9 respectively. The withdrawal of temporary measures disproportionately reduces incomes in less affluent counties, consistent with evidence that these payments played an important protective role for vulnerable and welfare-reliant households. For example, as shown in Figure 8, Donegal shows a relatively large decline in mean disposable income (around -1.53%) when temporary measures are removed, indicating greater exposure to the rollback of supports.

By contrast, the permanent measures alone display a broadly progressive geographic pattern. Freezing key income tax thresholds and bands generates fiscal drag as earnings rise, which weighs more heavily

Figure 9: Budget 2026, Change in disposable household income, permanent measures only

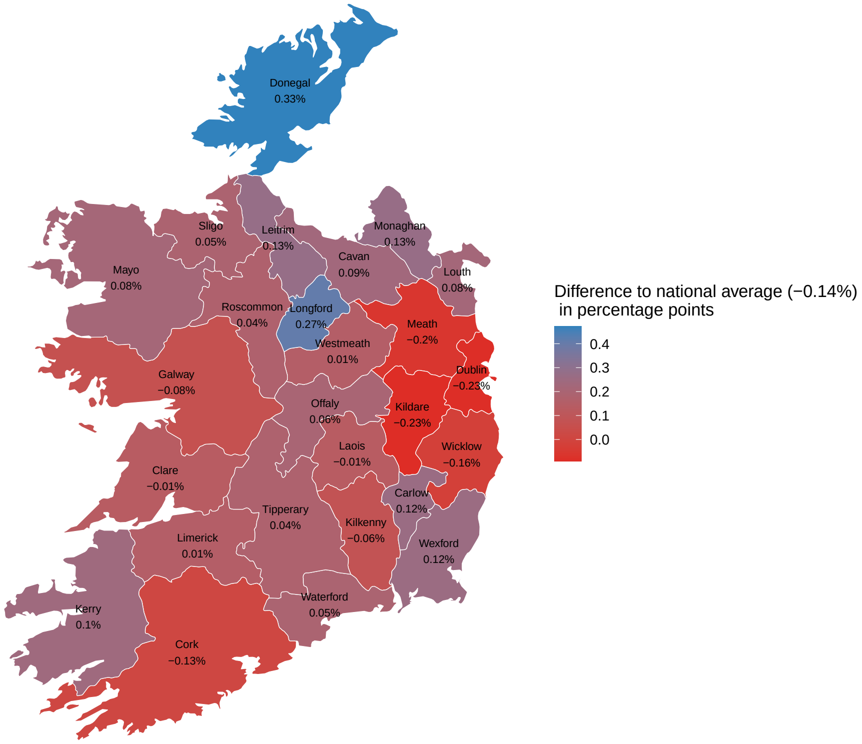

Notes. For each county, this Map reports the percentage change in mean disposable household income under the reform scenario (change in permanent measures only) relative to the price-uprated baseline scenario. Map shading shows the county's deviation from the national average (in percentage points: county change minus national change); blue indicates counties doing better than the national average and red indicates counties doing worse. Text labels show the percentage change between baseline and reform scenario. The national benchmark is computed using the original survey weight (*dwt*), while county estimates use the calibrated county weights derived from the deterministic method (averaged over bootstrapped replicate weights to account for sampling variability).

in counties with higher market incomes. At the same time, increases to core welfare rates and substantial increases to child-related supports benefit areas with higher benefit eligibility. Under permanent measures only, counties Donegal and Longford see the largest relative improvements (around +0.33% and +0.27% respectively) compared to the indexed baseline, while Dublin and much of the commuter belt tend to experience larger losses than the national average.

Taken together, the spatial pattern is consistent with the distributional results for the nationally representative sample reported by Bercholz and Simon [2025]: the permanent elements of Budget 2026 are progressive, but the withdrawal of temporary cost-of-living supports generates income losses across the distribution and is felt most strongly among lower-income households[21]. The county pattern therefore largely reflects differences in local income distributions and the extent to which households relied on the

---

[21] For point estimates with estimated confidence intervals, see Figure A.10 in the Appendix.

temporary supports that were withdrawn.

6 Conclusion

We provide a comprehensive comparison of two commonly used methods for Small Area Estimation in a unified data and policy setting. Conducting extensive internal validation of deterministic and probabilistic approaches, we show that the deterministic approach performs well at the county level under a rich set of constraints and is therefore suitable for integration with tax-benefit microsimulation. On this basis, we develop Geo-SWITCH, a spatially disaggregated extension of the SWITCH model that enables ex ante distributional analysis at the county level.

Our main methodological contribution is to bridge spatial microsimulation and tax–benefit modelling in a way that is operational for policy analysis. While spatial microsimulation methods typically focus on estimating distributions for single outcome variables, and tax-benefit models are only representative at the national level, we show that we can combine the two using calibration-based reweighting in a way that is fully compatible with an established microsimulation framework. A second contribution is to incorporate uncertainty quantification using bootstrap replicate weights allowing us to derive confidence intervals for spatially disaggregated policy outcomes.

In line with external data from the Irish Central Statistics Office, we find pronounced county-level differences in household disposable income. We incorporate uncertainty quantification using bootstrap confidence intervals, addressing a key gap in spatial microsimulation practice. Results demonstrate that Geo-SWITCH provides credible, policy-relevant estimates with moderate sampling variability.

Using Geo-SWITCH to study the distribution of free medical care in Ireland, we find substantial county-level variation, strongly linked to patterns in disposable income. Analysing the distributional impact of the most recent Irish budget, we find a similar level of regional differences. The withdrawal of temporary cost-of-living policies reduced the disposable income of those living in the North-West the most and those living in Dublin and surrounding areas the least. Looking at permanent changes to the tax and welfare system however, we find that these benefited the North-East of the country more than Dublin and surrounding areas.

Like other spatial microsimulation models, Geo-SWITCH relies on assumptions that affect what outcomes can be credibly estimated and the interpretation of results. The model produces county populations that are synthetic, hence using households from outside a county may introduce bias in the presence of specific local factors likely affecting the outcome. Furthermore, county calibration leads to more concentrated weights, meaning some county estimates may depend on relatively few donor households and are therefore less suitable for rare or weakly constrained outcomes. Finally, while bootstrap confidence intervals improve uncertainty quantification, they do not provide formal tests of differences between counties or fully capture bias from indirect estimates; further work could extend bootstrap methods to account jointly for bias and variance [Moretti and Whitworth, 2023].

With these caveats in mind, Geo-SWITCH is a powerful tool for understanding regional income disparities. Like many EU and OECD countries, Ireland has a nationally uniform tax and benefit system that is applied to household income that is unevenly distributed by region. This results in regional redistribution and tools such as Geo-SWITCH can be used to quantify the extent of this regional redistribution and support evidence-based decisions in taxation, and welfare.

Appendix A Supplementary Material

A.1 Level of observation and constraint variables

Figure A.1: Counties and NUTS2 areas in Ireland

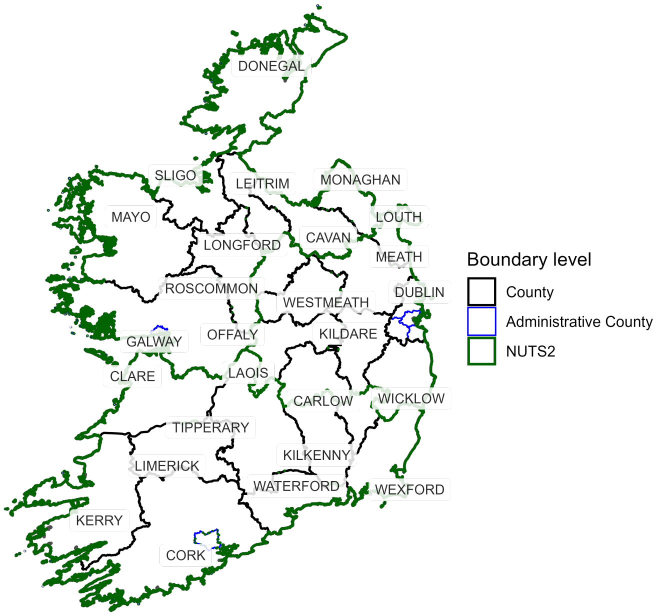

Notes. This Map shows the 26 counties of Ireland in black and the additional 5 counties in administrative county disaggregation in blue (Cork City, Galway City, Dún Laoghaire-Rathdown, Fingal and South Dublin) as well as the three NUTS2 areas used for the regional subsetting with green outline (North-West, South, East-Midland). *Source:* Tailte Éireann - Administrative Areas (2019).

Figure A.2: Number of households in SILC by administrative county

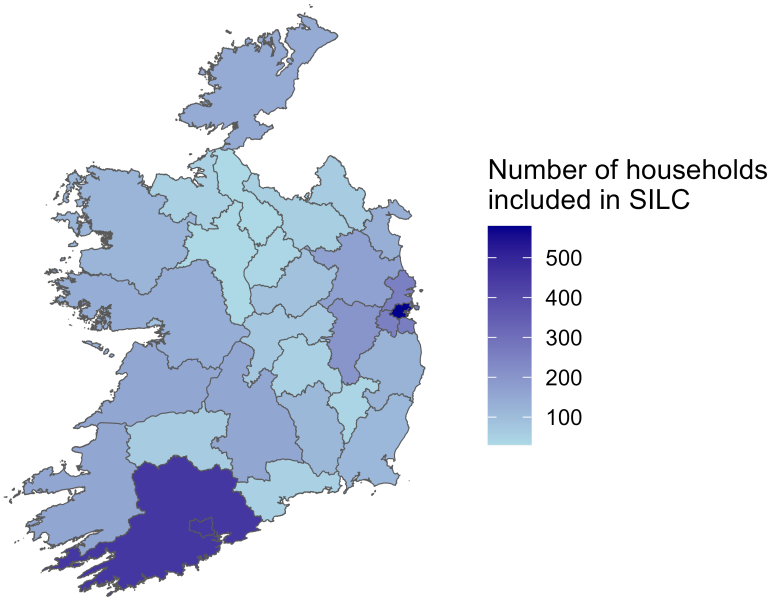

Notes. This Map shows the absolute number of household surveyed and included in the SILC survey data by administrative county (in 2022). *Source.* CSO, SILC RMF data (2022).

Table A.1: Overview of constraint variable selection and sources

Category	Variable(s)	Source	SILC [SWITCH var.]
Demographic	11 Age-sex bands *(0-4, 5-9, 10-14, 15-19, 20-29, ..., 70-79, 80+)*	Census (2022a)	*RB082, pb150, [dgn, dag]*
Education	5 levels (highest achieved: *In education, primary, lower secondary, upper secondary, third level*)	Highest level of education completed, Census (2022a) (F8038)	Combination of highest ISCED level attended: *PE041 [deh]* and current education status *PE010 [dec]*
Income	Gross household income: median and income brackets *(Administrative income data, including social welfare income)*	CSO Geographical Profile Admin. Income (2022b)	Calculated from market income, pensions and benefits: [*yem, yse, yot, ypp, ypr, poaps, poapu, poact, poals, pyr, psuwdct, pdi, poanc, bfalp, psuwdnc, bdinc, bsa00, bwkmt, bho00, bhout, bunnc, bsaot, bma, bdict, bunct, bhl, bfaot, bed, bch, buntr, yunsv, bca*]
Welfare	Welfare recipients: *Contributory pension, Illness benefit (contributory), One-Parent Family payment, Working Family [ayment; Carer's allowance; Jobseeker's assistance and benefit; State pension (Non-Contributory); Widow's pension (Contributory); Invalidity pension and Disability allowance.*	Department for Social Protection [2022]	Registry information: *[poact, bdict, bfalp, bwkmt, bca, bunnc, bunct, poanc, psuwdct, pdi, bdinc]*
Employment	Number employed, unemployed *(aged 15 and over, main status)*	Census (2022a), Population aged 15 and over (F7114)	*PL032 [les]*

Notes. All statistics are computed at county level (26 counties), variables are chosen in accordance with calibration variables in the national SWITCH model in Keane et al. [2023]. Education harmonisation: Census data contains residual categories ('Other', 'Not stated') that do not map directly to the five harmonised education levels in SILC. To ensure totals are consistent with the county population totals, we distribute the residual difference proportionally across the five education categories, preserving a county's observed education composition.

A.2 Detailed Methodological Setup

For the subsequent discussion, we formalise the problem as follows. The SILC survey comprises N households that are not spatially referenced. A set of J constraint variables is available both as microdata in the survey and as county-level totals from administrative sources. Let $T_{c,j}$ denote the total of a constraint j in county c, for example the total number of college-educated persons in Longford according to the 2022 Census. For each household $i = 1, ..., N$, the corresponding values of the constraint variables $j \in J$, aggregated from the individual level to the household level - are denoted by $\mathbf{x_i} = (x_{i,1}, ..., x_{i,J})$. In the college-education example, $x_{i,j}$ denotes the number of college-educated members in household i. Taking a set of initial survey weights at the household-level, $\mathbf{d} = (d_1, ..., d_N)$, our small-area estimation objective is, for each county c, to find a vector of adjusted weights $\mathbf{w_c} = (w_{1,c}, ..., w_{N,c})$ such that

$$T_{c,j} \approx \sum_{i=1}^{N} w_{i,c} x_{i,j}$$

ideally achieving equality. Note that weights may take the value zero, so that only households with $w_{i,c} > 0$ contribute to the synthetic population for county c, while those with $w_{i,c} = 0$ are effectively not selected for that county.

Deterministic Method

In our application, calibration converges to the set of household weights $\mathbf{w_c} = (w_{1,c}, ..., w_{N,c})$ where $\forall j \in J : T_{c,j} = \sum_{i=1}^{N} w_{i,c} x_{i,j}$ and which minimises the total distance $D = G(\mathbf{w_c}, \mathbf{d})$ from the starting weights under a chosen distance function G. As outlined in Deville and Särndal [1992], the solution weights are obtained numerically using Newton's method up to a researcher-specified precision, typically via statistical software; in our case, we use the *calibrate* function from the *survey* package in R. It is important to note that calibration is implemented independently for each county c, but, unlike the probabilistic method, but, unlike the probabilistic method, all population-level constraints are perfectly met once convergence is achieved.

The choice of distance function G has important implications for the set of solution weights. One choice is the chi-squared distance function $D = G(\mathbf{w_c}, \mathbf{d}) = \sum_{i=1}^{N} \frac{(w_c - d_i)^2}{2d_i}$. As this distance function does not impose any condition on what values the new weights can take, the solution set can contain very large, very small or even negative values. To avoid this, Deville et al. [1993] propose a different distance function which constrains the ratio of the new weight to the starting weight $\frac{w_c}{d_i}$ between a lower bound L and an upper bound U. As we want to avoid placing too much, or too little, weight on certain observations, we choose $U = 3$ and choose the smallest level of L for which the method converges. This distance function is the "logit" function in the *survey* package and is equivalent to the "ds" function in the *sreweight* Stata package which is used for calibration in the national SWITCH model.

Our choice of starting weights d is the *euroweight* variable provided in SILC. This weight has already been calibrated from the initial design weights by the CSO to attain national population totals for age-sex, NUTS3 region and household composition.

Probabilistic Method

The CMC method applies a rejection algorithm to a pool of sample data to achieve the county specific constraint totals $T_{c,j}$ outlined above. In a first step, the 4,660 households in our survey data are expanded by their survey weights $\mathbf{d}$ to give a starting pool P of households that is nationally representative and adjusted for the survey design. For a given county c, let $\overline{N}_c$ denote the set of households that have already been accepted. The current county-level total for constraint j is then accordingly $\overline{T}_{c,j} = \sum_{i\in\overline{N_c}} w_{i,c}x_{i,j}$, where the weight $w_{i,c}$ is equivalent to the number of times household i has been selected for county c.
A candidate household k is drawn from the pool and accepted if for all constraints j in J:

$$\overline{T}_{c,j} + x_{k,j} \leq T_{c,j}$$

If there exists at least one constraint j' in J such that

$$\overline{T}_{c,j'} + x_{k,j'} > T_{c,j'}$$

then candidate household k is rejected along with any other household k' where $x_{k',j'} \geq x_{k,j'}$.

Additionally, we take advantage of the flexibility of the CMC method to include distributional constraint variables, specifically the median gross household income. Let $T_{c,p50}$ denote the target median household income for county c, and let $\tau > 0$ be a tolerance parameter as allowed deviation.

A candidate household k that satisfies all other constraints $j \in J$, is then subjected to an additional check: out of the set of households previously accepted for a county, $\overline{N}_c$, it must hold that

$$\text{median}\{(x_{i,\text{inc}},\ldots,x_{\overline{N}_c,inc},x_{k,inc})\} \leq T_{c,p50} + \tau$$

for candidate household k to be accepted to $\overline{N_c}$. The algorithm continues in this way until all households i in the starting pool P have been either accepted or rejected.

Crucially, the final selection of households is made without replacement: once a household is assigned to county c, it cannot be reassigned. This differs from Simulated Annealing, which allows later state changes. This makes the procedure computationally efficient and feasible at scale, but it can impede convergence because misallocations cannot be corrected ex post.

A number of modifications are made to this basic underlying rejection algorithm to improve performance.

First, to reduce computational burden, we draw a working subsample $\overline{P} \subset P$ of candidate households and refresh this subsample from the full pool P once it is exhausted. In our application, a working sample size of 30,000 candidate households proved to be an effective compromise between computational efficiency and stability of the results. Second, at the initial stages of the algorithm, we restrict candidate draws to households with children to aid convergence towards the child-related constraint totals.

This rejection sampling continues until the current candidate set is exhausted. Once all candidate households for county c are processed, the median income is updated and compared to the target median $T_{c,p50}$. This discrepancy then guides the construction of a new candidate set drawn from N with income-based filtering: if the current median exceeds the target, subsequent draws are biased towards lower-income households; if it falls short, draws are biased towards higher-income households. At this stage, child-related constraints are also re-checked; if they remain unmet, the candidate pool is temporarily restricted to households with children until these constraints are satisfied.

To stop the algorithm cycling on a small set of duplicate households, we track what households are chosen on each refresh of the small sample pool $\overline{P}$. If the number of selected households is below a certain threshold and the same households are repeated then we relax the constraint set J in a deterministic schedule, allowing households to re-enter the set of feasible households.

A.3 Internal validation

Figure A.3: Internal validation, CMC method, county scatterplots by constraint

Notes. This Figure shows scatter plots comparing simulated totals from the CMC probabilistic method (y-axis) with target constraint totals (x-axis) for each county. Each point represents one county; red points indicate counties where the simulated value deviates by more than 10% from the corresponding target. The dashed 45-degree line denotes perfect agreement between simulated and target values. Each panel corresponds to a different constraint, with the panel title reporting the Pearson correlation coefficient between simulated and target totals across counties.
Sources: See Table A.1.

Table A.2: Census 2022 summary statistics by county

County	Household size share (%)												
	HH size 1	HH size 2	HH size 3	HH size 4	HH size 5+	Avg. # kids / HH	Third-level edu. (%)	Avg. age	Avg. HH social transfers (€)	Avg. HH market income (€)	Transfers share (%)	Market income share (%)	In work (%)
Carlow	8.2	20.3	20	24.6	26.8	1.36	39.3	38.8	13,227	52,703	20.1	79.9	54
Cavan	8.7	19.5	17.5	23.2	31.2	1.45	37.7	38.5	12,360	51,832	19.3	80.7	55.3
Clare	9.6	21.7	19	23.7	26	1.34	43.8	40.1	11,645	57,046	17	83	54.5
Cork	8.6	21.1	19.6	25.1	25.5	1.33	46.8	39.1	11,319	66,483	14.5	85.5	56.2
Donegal	10	21.5	18.3	22.7	27.4	1.39	36.8	40.1	13,700	40,559	25.2	74.8	51.1
Dublin	8.1	22.2	20.4	24.7	24.5	1.25	51.2	38.0	10,820	85,298	11.5	88.5	58.8
Galway	8.5	20.7	19.5	24.5	26.9	1.34	48.4	39.0	11,376	62,874	15.3	84.7	56
Kerry	10.8	23	19.1	23.1	24	1.29	41.3	41.5	12,392	48,815	20.2	79.8	52.9
Kildare	5.8	18.2	19.5	27.9	28.6	1.41	48.7	36.9	10,716	78,379	12	88	59.1
Kilkenny	8	20.5	19.6	24.4	27.4	1.37	43.8	39.7	12,175	60,493	16.8	83.2	55.6
Laois	7.3	18.1	18.7	25.7	30.2	1.48	39.3	37.3	12,328	59,310	17.2	82.8	55.9
Leitrim	11.8	22.5	17.6	20.9	27.1	1.35	42.2	40.7	12,457	45,515	21.5	78.5	52.8
Limerick	8.9	22	20	23.5	25.6	1.32	41.2	39.3	12,782	59,271	17.7	82.3	53.3
Longford	9.6	20.4	19.6	22	28.4	1.41	34.2	38.8	13,723	46,079	22.9	77.1	53.2
Louth	8.1	19.2	20.1	25.1	27.5	1.42	39.6	38.2	13,384	55,047	19.6	80.4	53.5
Mayo	10.8	23.1	18	21.7	26.5	1.32	39.2	41.6	12,735	48,880	20.7	79.3	52.3
Meath	5.7	17.2	18.7	28	30.4	1.46	45.2	36.9	11,275	74,408	13.2	86.8	59.3
Monaghan	8.7	19.3	17.3	22.6	32	1.49	36.6	38.7	12,443	49,891	20	80	56.1
Offaly	7.8	19.3	18.7	24.7	29.4	1.41	37	39.0	13,528	54,909	19.8	80.2	53.4
Roscommon	10	21.5	18.7	22.9	26.8	1.34	40.3	40.8	12,539	51,536	19.6	80.4	52.7
Sligo	11.2	24.1	18.7	21.8	24.2	1.30	44.5	40.7	12,090	52,247	18.8	81.2	53.4
Tipperary	9.6	21.5	19.3	23.3	26.3	1.36	38.4	40.3	13,025	53,881	19.5	80.5	54.4
Waterford	9.7	23.1	19.7	23.6	23.9	1.30	41	39.9	12,498	54,565	18.6	81.4	53.8
Westmeath	8.3	20.2	19.4	24.3	27.7	1.38	41.7	38.6	12,410	58,409	17.5	82.5	55
Wexford	8.8	21.4	19.1	24.1	26.5	1.32	37	40.0	13,178	50,666	20.6	79.4	52.9
Wicklow	7.2	19.9	19.5	26.6	26.7	1.36	47.7	39.1	11,822	71,645	14.2	85.8	55.8

Notes: This Table reports county-level summary statistics on demographic characteristics, income, employment and transfers, drawn from the 2022 Census [Central Statistics Office, 2022a].

Table A.4: Convergence rates by county, deterministic method

County	Convergence rate (in % of 1,000 replications)
Offaly	96.4
Louth	96.7
Carlow	97.1
Kilkenny	97.7
Tipperary	97.7
Wicklow	97.7
Dublin	97.9
Roscommon	97.9
Wexford	98.3
Meath	98.6
Leitrim	98.7
Longford	98.8
Monaghan	98.8
Waterford	98.8
Westmeath	98.8
Donegal	99.3
Limerick	99.3
Sligo	99.3
Kerry	99.5
Cork	99.6
Cavan	99.9
Kildare	99.9
Laois	99.9
Clare	100.0
Galway	100.0
Mayo	100.0

Notes: This Table reports convergence rates per 1,000 bootstrap replications of the deterministic reweighting (using logit distance function) by county in ascending order.

Table A.5: Internal validation, constraint fit of all variables, CMC method

	Total					Correlation coef.	
Constraint	target	CMC	Bias (units)	Rel. bias	RMSE (units)	Pearson	Spearman
Employment							
Employed	2,320,297.0	2,320,297.0	0.0	0.0%	0.00	1.000	1.000
Unemployed	210,802.0	210,802.0	0.0	0.0%	0.00	1.000	1.000
NILF	2,618,040.0	2,618,040.0	0.0	0.0%	0.00	1.000	1.000
Age-sex brackets							
0_4_M	151,408.0	151,476.6	68.6	0.0%	13.38	1.000	1.000
5_9_M	175,470.0	175,619.2	149.2	0.1%	29.05	1.000	1.000
10_14_M	191,114.0	191,280.7	166.7	0.1%	32.46	1.000	1.000
15_19_M	172,342.0	172,435.5	93.5	0.1%	18.26	1.000	1.000
20_29_M	302,844.0	286,309.3	-16,534.7	-5.5%	2,065.68	0.999	0.994
30_39_M	343,729.0	343,491.9	-237.1	-0.1%	321.41	1.000	0.999
40_49_M	385,735.0	388,294.8	2,559.8	0.7%	414.17	1.000	1.000
50_59_M	321,892.0	330,740.6	8,848.6	2.7%	953.59	1.000	0.999
60_69_M	251,813.0	265,039.8	13,226.8	5.3%	1,442.89	0.999	0.998
70_79_M	173,007.0	169,241.0	-3,766.0	-2.2%	342.44	0.999	0.984
80+_M	75,195.0	73,527.7	-1,667.3	-2.2%	161.26	0.999	0.990
0_4_F	144,007.0	144,064.8	57.8	0.0%	11.28	1.000	1.000
5_9_F	167,200.0	167,285.7	85.7	0.1%	16.69	1.000	1.000
10_14_F	183,088.0	183,228.6	140.6	0.1%	27.38	1.000	1.000
15_19_F	165,286.0	165,362.9	76.9	0.0%	14.98	1.000	1.000
20_29_F	300,107.0	275,026.5	-25,080.5	-8.4%	4,437.29	0.997	0.999
30_39_F	371,363.0	377,420.3	6,057.3	1.6%	657.92	1.000	0.999
40_49_F	399,293.0	401,910.0	2,617.0	0.7%	339.53	1.000	1.000
50_59_F	325,276.0	338,552.8	13,276.8	4.1%	1,416.32	1.000	0.999
60_69_F	259,001.0	267,360.9	8,359.9	3.2%	936.77	1.000	0.997
70_79_F	184,137.0	184,646.1	509.1	0.3%	212.20	1.000	0.999
80+_F	105,832.0	96,823.3	-9,008.7	-8.5%	546.49	0.997	0.964
Education							
In education	1,634,624.0	1,610,620.7	-24,003.3	-1.5%	4,780.64	1.000	0.999
Primary	368,778.0	395,598.4	26,820.4	7.3%	2,023.57	0.997	0.990
Lower-secondary	494,331.0	466,293.1	-28,037.9	-5.7%	2,134.66	0.996	0.989
Upper-secondary	962,367.0	828,068.8	-134,298.2	-14.0%	8,840.47	0.999	0.992
Third-level	1,689,039.0	1,848,557.9	159,518.9	9.4%	11,888.24	1.000	0.999
Benefit receipt							
State pension (contr.)	437,688.0	435,128.2	-2,559.8	-0.6%	2,063.33	0.999	0.992
State pension (non-contr.)	97,703.0	66,087.3	-31,615.7	-32.4%	1,686.05	0.934	0.904
Widower's pension	117,238.0	117,543.1	305.1	0.3%	963.53	0.993	0.968
Invalidity pension	54,685.0	56,369.7	1,684.7	3.1%	355.33	0.997	0.985
One-Parent Family payment	43,202.0	39,389.6	-3,812.4	-8.8%	747.81	0.994	1.000
Working Family payment	46,968.0	48,624.9	1,656.9	3.5%	258.76	0.999	0.996
Jobseeker's allowance	146,033.0	167,042.6	21,009.6	14.4%	2,436.61	0.993	0.961
Jobseeker's benefit	40,790.0	43,497.9	2,707.9	6.6%	431.99	0.997	0.975
Carer's allowance	92,139.0	97,533.3	5,394.3	5.9%	612.63	0.999	0.999
Illness benefit	58,912.0	62,209.2	3,297.2	5.6%	202.94	1.000	0.997
Disability allowance	157,738.0	172,246.4	14,508.4	9.2%	1,640.66	0.996	0.984
Median income	55,413.8	53,128.0	-2,285.9	-4.1%	2,760.33	0.987	0.980
1-19,999	232,402.5	320,892.9	88,490.4	38.1%	7,739.34	0.984	0.973
20,000-39,999	363,663.4	384,248.1	20,584.7	5.7%	1,021.57	1.000	0.996
40,000-59,999	321,553.8	350,772.0	29,218.2	9.1%	1,331.43	1.000	0.984
60,000-79,999	260,030.7	278,506.6	18,475.9	7.1%	1,154.82	1.000	0.999
80,000-99,999	195,322.7	201,008.8	5,686.1	2.9%	309.26	1.000	0.995
100,000-119,999	140,380.9	137,760.2	-2,620.6	-1.9%	373.17	0.999	0.983
120,000-139,999	98,297.6	98,685.1	387.5	0.4%	126.50	1.000	0.994
140,000-159,999	66,633.0	67,677.2	1,044.3	1.6%	51.83	1.000	0.996
160,000-179,999	44,965.9	42,836.4	-2,129.5	-4.7%	401.49	1.000	0.997
180,000-199,999	29,796.2	24,179.7	-5,616.5	-18.8%	675.81	0.996	0.970
200,000+	83,681.3	74,539.7	-9,141.6	-10.9%	2,012.05	0.997	0.997

Notes. This Table summarises model fit of the CMC method in achieving the target constraints. The ”target” values are drawn from the Census 2022 and other administrative sources. Note that for education the numbers for categories ”Other” and ”No response” were proportionately distributed to all other categories by county.

Sources: See Table A.1.

A.4 Constraint relevance

Regression analysis We assess the relevance of constraints imposed in the small-area estimation using survey-weighted regressions of (log) disposable income at the household level which is the primary outcome variable of the SWITCH microsimulation model.

Table A.6 summarises the results, using the constraint variables (adjusted as indicator variables at household level) as predictors of log household disposable income. The constraint set shows strong relevance - taken together, the age structure, education, employment and welfare indicators explain around 46% of the variation in household income ($R^2 = 0.46$). Households with higher education (Third level) and more employed members have significantly higher incomes, while a greater number of unemployed members and receipt of means-tested benefits (e.g. unemployment and working-age payments) are associated with lower incomes, as expected. Lower education categories are negatively related to income relative to the omitted group (In education), and several age-bracket counts are strongly associated with income, indicating that the demographic composition of a household matters for household disposable income. Overall, the results support the use of these variables as constraints in the small-area estimation, as they capture key gradients in disposable income.

Random Forest prediction As a non-parametric test of the relevance of our constraint variables, we estimate a survey-weighted random forest model of (log) disposable income at the individual level, allowing all available SWITCH output variables to enter as potential predictors. The random forest grows many decision trees on bootstrap samples, each time considering only a random subset of covariates at each split; predictions are obtained by averaging across trees. This flexible setup captures non-linearities and interactions without requiring to specify a particular functional form beforehand.

The random forest results in Table A.7 confirm that our chosen constraints are relevant for explaining disposable income: the most important predictors are components of gross household income and closely related labour-market and welfare variables (e.g. employment income, self-employment income, investment income, pensions, and key benefits), which directly map onto our constraint set. Characteristics such as disability status or housing tenure also emerge as important; these are not used as calibration constraints but instead serve as external validation variables to assess the distributional fit of the small-area estimates.

Table A.6: Constraint relevance, regression results

	Dep. var.: Disposable HH income (log)
Intercept	6.979***
	(0.029)
Highest education level achieved	
Primary	-0.140***
	(0.030)
Lower secondary	-0.117***
	(0.028)
Upper secondary	-0.028
	(0.025)
Third level	0.087***
	(0.024)
Employment status	
Employed	0.422***
	(0.029)
Unemployed	-0.319***
	(0.041)
Benefit receipt	
State pension contr. (poact)	0.040
	(0.035)
State pension non-contr. (poanc)	-0.249***
	(0.050)
Illness benefit (bdict)	-0.078.
	(0.047)
One-parent family payment (bfalp)	-0.095
	(0.062)
Basic Suppl. Welfare Allowance (bsa00)	-0.308.
	(0.165)
Working Family payment (bwkmt)	-0.491***
	(0.051)
Carer's Allowance (bca)	-0.014
	(0.040)
Jobseeker's Allowance (bunnc)	-0.320***
	(0.042)
Jobseeker's Benefit (bunct)	-0.127.
	(0.065)
Widow/Widower's/Surviving Partners' benefit (psuwdct)	0.003
	(0.040)
Invalidity pension (pdi)	-0.111*
	(0.050)
Disability allowance (bdinc)	-0.229***
	(0.034)
Observations	4,642
Adj. R^2	0.459

Notes: Survey-weighted OLS. Robust standard errors in parentheses. Household-level dummies indicate whether *any* member is unemployed, employed, or in receipt of the corresponding payment. Omitted categories are "In education" and "Not in the labour force" for education and employment status. Age-composition controls (household counts in 11 age brackets) are included in the regression but omitted from the table; estimates are all positive and jointly significant ($p < 0.001$), confirming that age structure is strongly associated with income. Significance levels: . $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table A.7: Constraint relevance, random forest prediction

Rank	Importance	Type	Description (variable)
1	2.7500	Constraint (HH income)	Employee income (`yem`)
2	0.2468		Hours worked (`lhw`)
3	0.1852	Constraint-related	Work history / months in work (`liwwh`)
4	0.1183	Constraint	State pension (contributory) (`poact`)
5	0.1085	External validation	Disability indicator (`ddi`)
6	0.0923		Labour market: industry (`lindi`)
7	0.0874	Constraint-related (age)	Pensioner / months retired (`lpemy`)
8	0.0721		Family type (`famtype`)
9	0.0590	Constraint (HH income)	Income from investment (`yiy`)
10	0.0584	Constraint (HH income)	Financial assets / capital (`afc`)
11	0.0560	Constraint-related	Education in years (`dey`)
12	0.0405	Constraint	In education indicator (`eduIn.edu`)
13	0.0355		Housing cost / expenditure (`xhc`)
14	0.0346		Previous employment income (`yempv`)
15	0.0340		Mortgage payment (`xhcmomi`)
16	0.0260		Occupational indicator (`loc`)
17	0.0226		Market value of main residence (`amrmv`)
18	0.0196	Constraint	Disability benefit (`bdinc`)
19	0.0170	Constraint (HH income)	Other income sources (`yot`)
20	0.0107	Constraint (HH income)	Self-employment income (monthly) (`ysemy`)
21	0.0104	Constraint (HH income)	Property income (`ypr`)
22	0.0101	External validation	Main residence tenure type (`amrtn`)
23	0.0098	Constraint	In work (labour market status) (`liwftmy`)
24	0.0097	Constraint (HH income)	Initial self-employment income (`yse_0a`)
25	0.0097	Constraint (HH income)	Self-employment income (`yse`)
26	0.0073	Constraint	Widow/er or surviving partner benefit (contrib.) (`psuwdct`)
27	0.0072	Constraint	Carer's allowance (`bca`)
28	0.0064		Private pension contribution (`xpp`)
29	0.0062		Property value (`amrrm`)

Notes: Importance values are permutation-based variable importances from the random forest model. "Constraint" indicates variables used as constraints in the small-area estimation; "External validation" marks variables compared to external statistics but not constrained on. Labels and descriptions should be aligned with the SWITCH/SILC documentation.

A.5 Regional Subsetting

We also consider a regional subsetting strategy, which utilises the observed location of survey households to restrict the reweighting procedure to households from the three NUTS2 regions in Ireland (Eastern and Midland, Southern, and Northern and Western, see Map A.1). Let $r \in \{1,2,3\}$ index the three regions, then we know where each household in the survey data is located in one of the regions $r(i)$, and each target county is located in a region $r(c)$.

In the regional subsetting specification, the synthetic population for county c is then constructed only from households located in this region, hence we impose for the weights $w_{i,c} = 0$ whenever $r(i) \neq r(c)$. The calibration problem is otherwise unchanged (see section 4.1). This approach addresses the restrictive spatial homogeneity assumption in that we allow for regional factors at the NUTS2 level to affect household characteristics that are not captured by the constraint set. However, this comes with a trade-off. By limiting the households to those located in the corresponding region of a county, the effective sample size is reduced which can lead to more concentrated weight distributions, with some households receiving relatively large weights. While, in principle, we could further restrict the set of households to those in the target county as we have this information (see Figure A.2, the resulting small cell sizes would make it generally unfeasible to satisfy the full set of constraints, so the NUTS2-level restriction provides a practical compromise.

However, inspection of the weight distributions (see Figure 1) shows that regional subsetting leads to very high weight concentration. This is further reflected in the confidence interval estimates using the bootstrapping procedure described in section 4.4; as shown in Figure A.4, the intervals are noticeably wider and frequently overlap across counties, consistent with the greater sampling uncertainty induced by relying on fewer households. This is especially the case for households located in the NUTS region ”North-West”, e.g. Donegal, which has a smaller sub-sample of households interviewed for SILC (see Map A.2) and therefore uncertainty is much larger. In contrast, for counties in the NUTS region East-Midlands (e.g. Dublin, Meath, Kildare), the point estimates and confidence intervals are much more aligned.

Thus, we conclude that drawing from the full national household sample is necessary to produce stable, representative county-level estimates in our setting, even if this might introduce bias which should be further investigated.

Figure A.4: Median household disposable income (2022) with 95% CI, regional subsetting

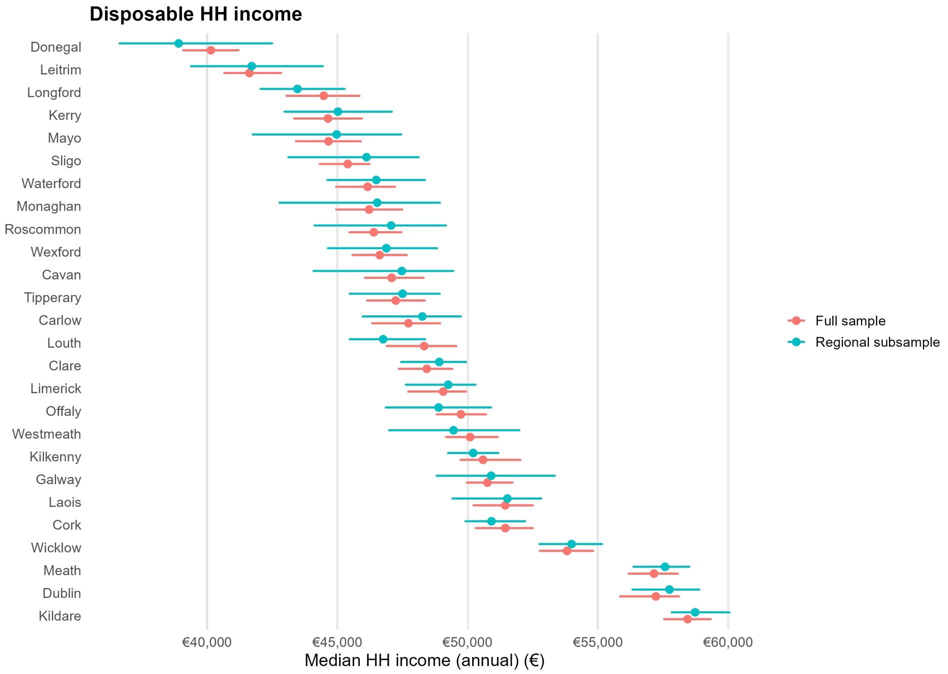

Points = Rao & Wu bootstrap means; lines = 95% CI

Notes. Red points show county-level estimates from the full sample, and blue points show the corresponding calibration estimates for the regional subsample; horizontal lines indicate 95% Rao and Wu bootstrap estimated confidence intervals [Rao and Wu, 1993]. Counties are ordered by increasing disposable household income.

A.6 External Validation

Figure A.5: External validation, heatmap of bias, secondary constraint set

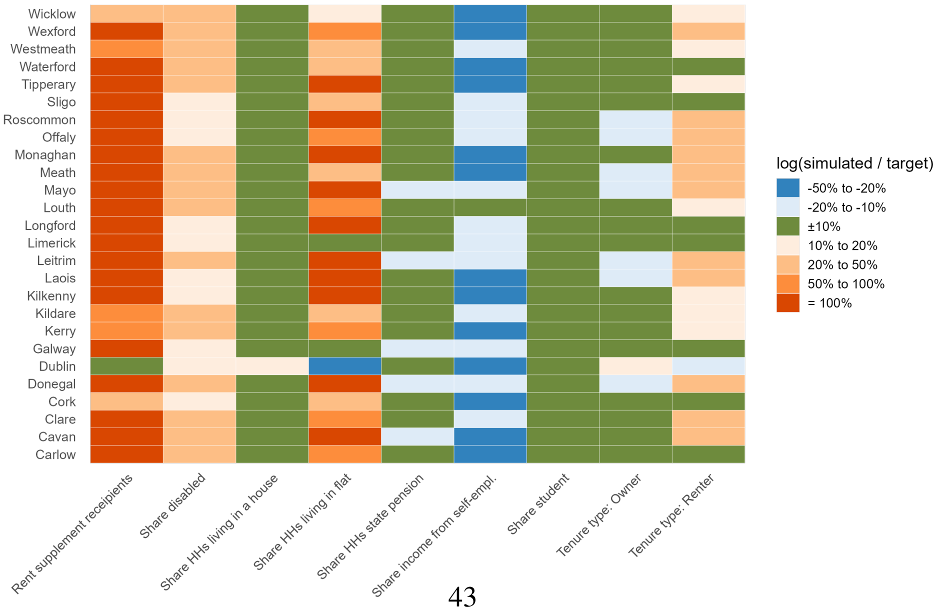

Notes. This heatmap shows correspondence as expressed by log ratio bias between small-area estimates (calibration, logit distance function) and the secondary set of constraints (as columns) by county (rows). *Sources:* See Table 2.

Table A.8: External validation, alternative calibration including household size as constraint

	Mean					Correlation coef.	
Variable set	target	simulated	Bias (units)	Rel. bias	RMSE (units)	Pearson	Spearman
Primary							
Share HHs majority benefits	0.303	0.300	-0.004	-0.982%	0.008	0.990	0.981
Share HHs state pension	0.151	0.141	-0.010	-5.978%	0.014	0.930	0.921
Share retired	0.167	0.160	-0.008	-4.380%	0.010	0.966	0.948
Share students	0.108	0.144	0.037	33.821%	0.037	0.864	0.900
Child benefit recip. (abs.)	24,744	24,934	190	-0.149%	1,263	0.999	0.993
Rent supplement recip. (abs.)	348	320	-28	164.112%	583	0.933	0.651
Share in HH size: 1	0.089	0.089	0.000	0.000%	0.000	1.000	1.000
Share in HH size: 2	0.210	0.210	0.000	0.000%	0.000	1.000	1.000
Share in HH size: 3	0.192	0.192	0.000	0.000%	0.000	1.000	1.000
Share in HH size: 4	0.242	0.242	0.000	0.000%	0.000	1.000	1.000
Share in HH size: 5+	0.268	0.268	0.000	0.000%	0.000	1.000	1.000
HH disp. income (total)	5,194,834,615	4,214,272,993	-980,561,622	-19.172%	1,754,817,981	0.999	0.993
HH disp. income per capita	24,220	19,715	-4,504	-18.396%	4,750	0.639	0.471
HH gross median income	55,414	55,964	550	0.982%	776	0.996	0.999
HH gross income (total)	7,681,454,615	5,754,409,151	-1,927,045,464	-25.860%	3,593,325,212	0.999	0.996
Secondary							
Share HHs living in a house	0.922	0.879	-0.043	-4.455%	0.058	0.819	0.582
Share HHs living in flat	0.078	0.121	0.043	83.222%	0.058	0.819	0.582
Share income from self-empl.	0.114	0.094	-0.020	-16.592%	0.023	0.414	0.351
Share disabled	0.050	0.055	0.005	10.615%	0.006	0.957	0.936
Tenure type: Owner	0.708	0.677	-0.030	-3.967%	0.051	0.453	0.426
Tenure type: Renter	0.292	0.323	0.030	13.116%	0.051	0.453	0.426

Notes. This Table summarises external validation results at the county level for the constraint set including household size. All statistics are averaged across counties, giving each county equal weight (i.e. results are not population-weighted). The Table therefore reflects the average model fit across counties rather than the fit for the average individual. Shares of retired persons, students, and disabled individuals are calculated relative to the population aged 15 and over. Child benefit and rent supplement are reported as absolute numbers of recipients. Household size shares refer to the proportion of the population living in households of size 1, 2, 3, 4, or 5+. Here, they are used as internal constraints so hit all the benchmarks perfectly. Simulated county-level means are additionally averaged across bootstrap replicate weights, so reported estimates reflect the mean across bootstrap iterations and are thus robust to sampling variability.

Figure A.6: External validation, heatmap of percentage error, primary constraint set

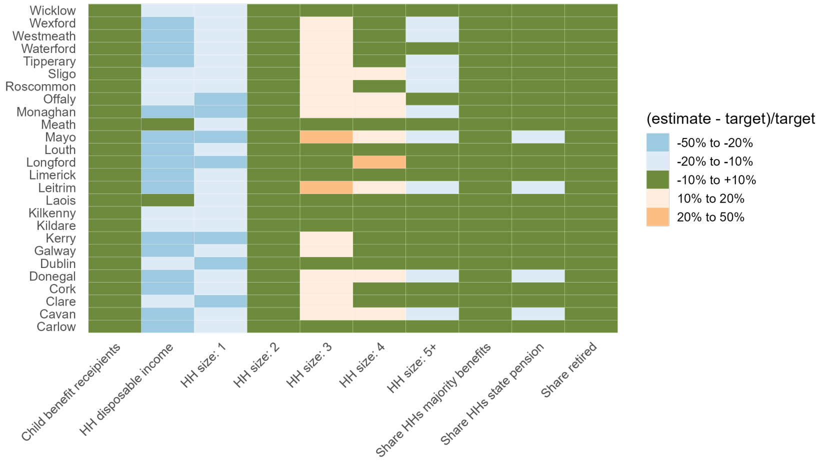

Notes. This heatmap shows correspondence expressed as percentage error between small-area estimates (calibration, logit distance function) and the primary set of constraints (as columns) by county (rows).
Sources: See Table 2.

Table A.9: External validation: simulated / target

				HH size														
County	Rent supplement	Child benefit	Disp. income	1	2	3	4	5+	Share HHs in flat	Share HHs in house	Share HHs majority benefits	Share HHs state pension	Share income self-employ.	Share retired	Share student	Share disabled	Tenure type: Owner	Tenure type: Renter
National	177.7	104.2	82.7	94.6	97.2	99.8	107.2	97.4	99.4	100.1	97.9	100.7	86.5	91.9	128.0	108.6	100.0	100.0
Carlow	241.6	98.5	69.2	87.4	98.7	109.4	107.2	91.3	167.8	94.7	97.3	103.4	76.7	103.3	99.9	120.8	95.3	109.9
Cavan	232.5	103.5	76.5	82.0	98.0	117.3	115.8	84.8	227.8	93.3	98.1	89.3	74.7	103.0	98.1	120.0	91.4	122.5
Clare	760.7	103.2	83.8	78.0	97.7	116.0	107.4	91.6	184.7	94.4	103.1	91.9	82.8	100.7	100.9	123.3	91.3	126.7
Cork	125.7	105.5	77.1	85.3	98.7	111.3	103.9	93.5	125.5	97.2	101.5	96.8	75.4	107.1	102.1	118.8	100.5	99.0
Donegal	1341.8	106.2	77.4	84.2	97.6	114.6	112.0	88.0	252.9	92.8	94.2	80.1	83.5	90.6	98.8	127.9	86.0	138.1
Dublin	94.3	107.2	80.1	78.1	96.7	109.1	108.3	94.2	52.7	117.6	99.3	104.8	79.6	101.1	103.6	119.9	110.5	83.3
Galway	244.8	107.0	76.6	81.9	96.7	111.2	108.8	92.2	107.7	99.0	101.4	90.4	84.5	103.4	103.3	116.0	98.5	103.5
Kerry	181.3	102.2	74.6	74.2	98.0	118.2	107.7	91.5	191.8	94.2	102.0	91.0	62.3	99.0	97.6	122.9	92.8	119.3
Kildare	156.5	100.6	88.3	89.4	97.3	105.3	103.4	96.9	122.6	97.5	98.9	98.3	87.3	106.1	101.7	122.4	94.8	114.8
Kilkenny	680.0	105.4	84.3	82.1	99.4	107.8	107.6	93.4	217.7	93.7	98.5	93.4	67.3	102.5	98.8	115.9	94.1	116.9
Laois	697.6	104.8	90.2	81.8	97.4	109.0	109.8	92.1	222.8	93.2	95.9	94.6	78.4	108.4	97.3	118.9	90.6	126.4
Leitrim	731.6	103.2	77.4	84.2	99.6	120.5	115.0	82.2	246.4	92.4	99.0	89.4	90.8	94.9	96.6	126.6	89.9	128.1
Limerick	347.3	106.9	75.1	87.0	95.7	108.4	109.9	92.5	108.6	98.8	97.9	97.6	81.3	99.6	105.1	114.1	96.5	107.9
Longford	548.5	107.1	78.0	67.5	91.8	106.6	122.4	95.0	201.8	93.9	92.8	92.5	81.6	98.1	97.8	118.3	98.2	103.4
Louth	559.3	103.5	78.9	84.7	103.3	104.5	108.1	91.5	162.1	95.1	94.5	99.0	96.7	99.4	99.9	121.7	92.9	116.3
Mayo	369.5	105.6	78.0	79.3	100.6	120.5	113.2	83.2	211.0	93.8	98.9	85.2	83.5	98.0	97.0	121.3	90.3	130.0
Meath	487.3	99.6	91.9	86.0	98.8	107.5	103.3	95.7	138.5	96.5	98.1	94.2	79.8	106.6	99.3	124.5	90.1	135.0
Monaghan	257.6	105.3	74.0	79.4	98.7	115.7	119.1	84.3	222.2	93.6	99.6	94.0	66.1	100.5	97.6	123.5	91.4	123.4
Offaly	518.4	105.2	85.1	78.9	99.3	110.9	110.8	90.0	198.6	94.4	96.3	95.2	80.9	102.0	98.0	119.5	90.2	129.0
Roscommon	391.1	104.5	83.3	82.1	102.4	111.6	109.1	88.8	323.1	92.0	98.7	92.9	85.1	101.3	98.4	118.3	89.9	134.1
Sligo	416.4	107.8	81.4	82.4	95.0	117.6	110.8	89.8	147.5	95.7	99.2	93.6	86.1	95.1	103.6	115.1	96.6	108.0
Tipperary	542.0	103.9	73.1	86.1	100.9	110.7	108.8	88.7	250.2	92.7	97.7	96.4	68.8	99.5	98.6	120.1	94.0	115.6
Waterford	908.0	106.1	76.8	88.8	95.9	110.9	105.7	93.9	143.7	95.8	95.1	94.6	79.1	97.8	100.5	123.0	97.6	105.0
Westmeath	171.6	101.4	79.1	81.8	98.5	111.6	109.8	89.8	122.3	97.6	99.5	103.4	80.8	101.1	98.5	126.3	94.1	114.3
Wexford	284.7	102.7	74.6	82.2	98.1	112.8	109.2	89.8	198.7	94.3	98.3	95.6	68.7	99.0	97.0	124.0	91.3	123.0
Wicklow	122.3	100.3	85.6	83.5	100.5	107.7	102.0	96.4	111.7	98.7	99.7	99.1	75.4	105.8	99.4	123.9	95.8	111.8

Figure A.7: External validation, heatmap of percentage error, secondary constraint set

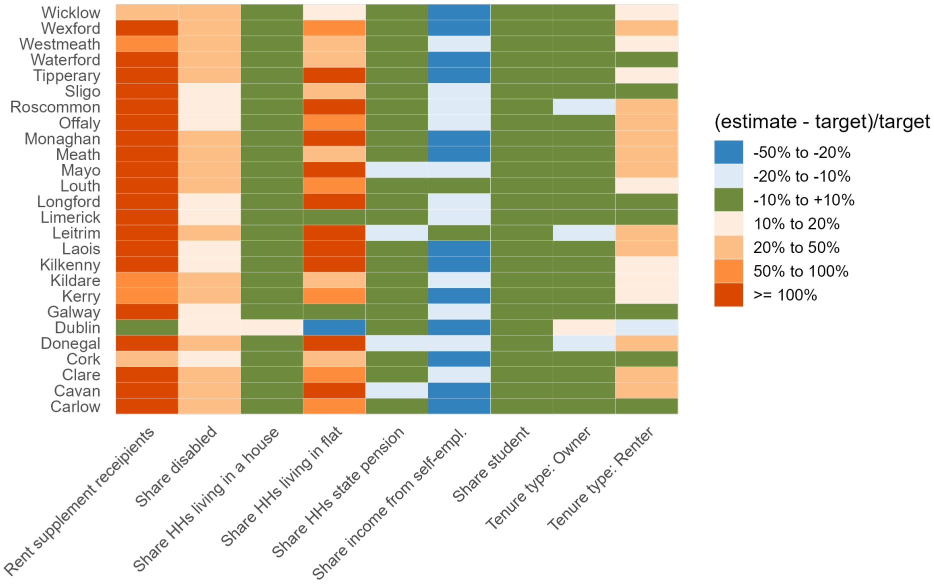

Notes. This heatmap shows correspondence expressed as percentage error between small-area estimates (calibration, logit distance function) and the secondary set of constraints (as columns) by county (rows).
Sources: See Table 2.

A.7 Applications

Figure A.8: Market household income distribution (2022)

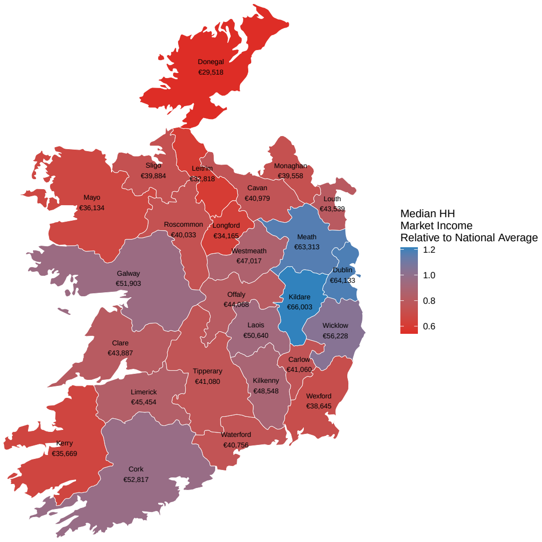

(a) Market income

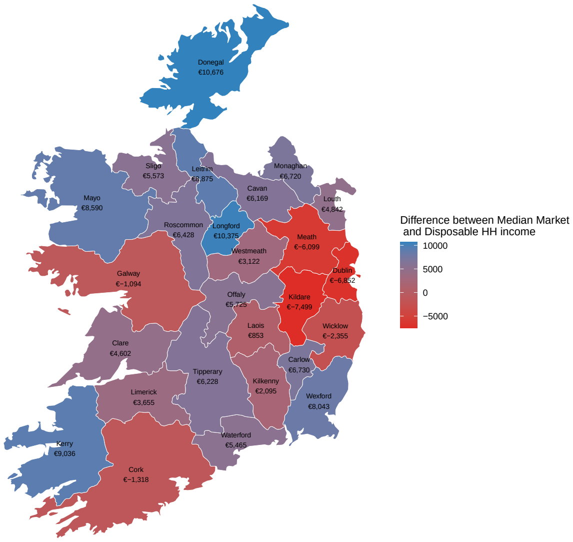

(b) Difference between disposable and market income

Notes. This Map shows annual median household market income (income from employment, self-employment, investments) as label. Colours represent each county's value relative to the national benchmark (normalised to 1.0) from the standard (non-georeferenced) microsimulation model. The legend is scaled to the observed range of county-to-national ratios, so 1.0 is not necessarily centered on the colour bar. This is the mean value across bootstrapped replications of the deterministic method to ensure robustness to sample variability.

Figure A.9: Medical Card, share individuals eligible (2022) with 95% CI

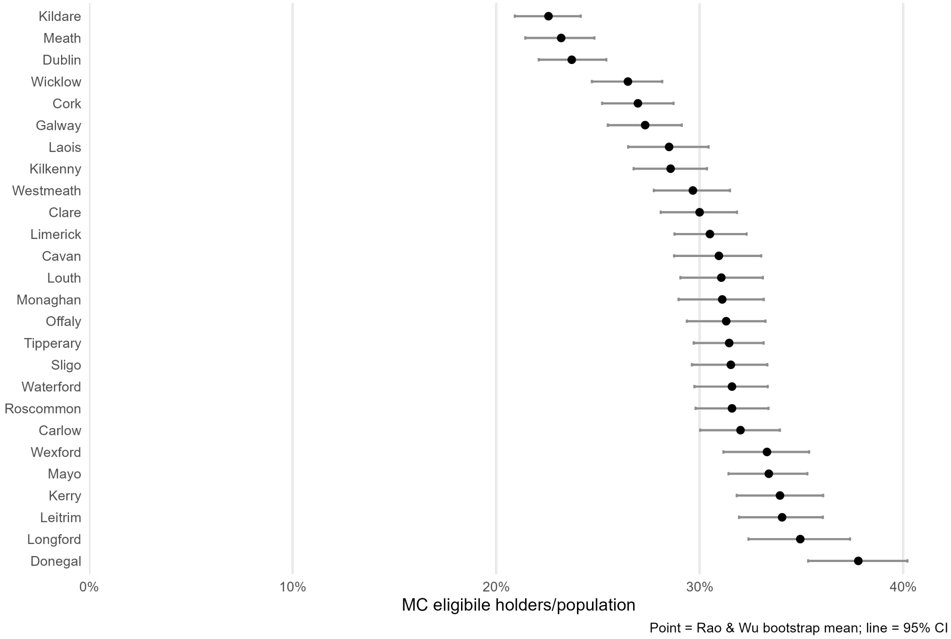

Notes. Points show county-level estimates based on the full sample using bootstrapping; horizontal lines indicate 95% bootstrap confidence intervals [Rao and Wu, 1993]. Counties are ordered by increasing point estimate for share of individuals in a county eligible for a Medical Card.

Figure A.10: Budget 2026, Change in disposable household income with 95% CI

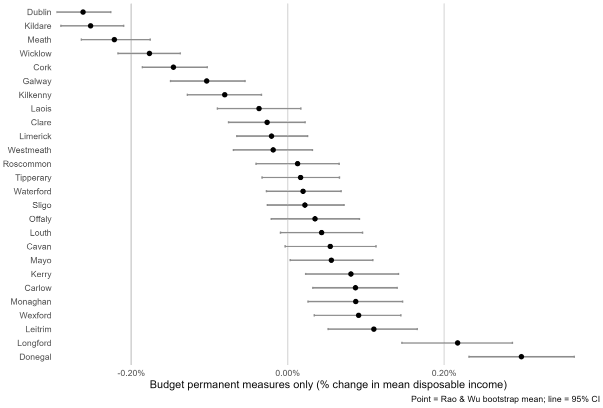

(a) Permanent changes only

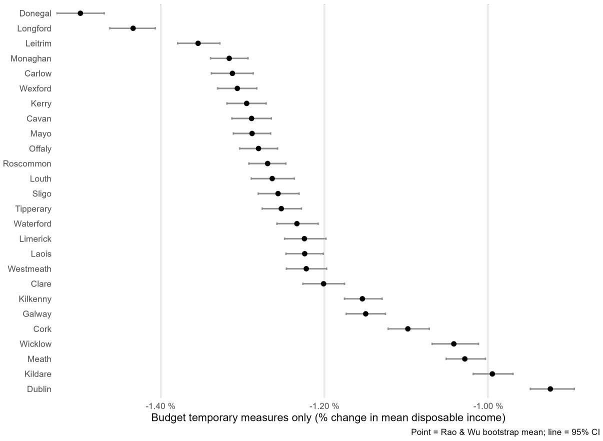

(b) Temporary changes only

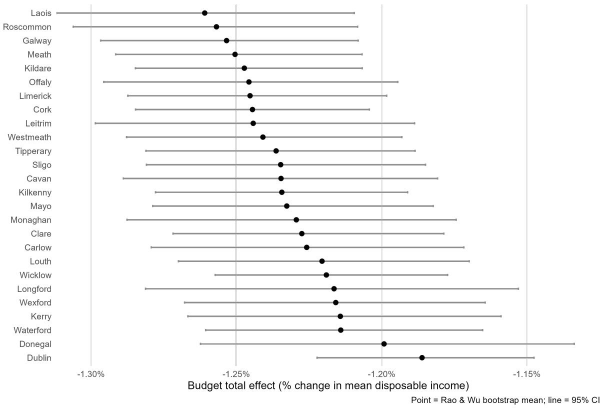

(c) Total change

Notes. Points show county-level estimates based on the full sample using bootstrapping; horizontal lines indicate 95% bootstrap confidence intervals [Rao and Wu, 1993]. Counties are ordered by point estimate for change in disposable household income.

Disclaimer

Results are based on analysis of strictly controlled Research Microdata Files provided by the Central Statistics Office (CSO). The CSO does not take any responsibility for the views expressed or the outputs generated from this research.

References

Dimitris Ballas, Graham Clarke, Danny Dorling, Heather Eyre, Bethan Thomas, and David Rossiter. Simbritain: A Spatial Microsimulation Approach to Population Dynamics. *Population, Space and Place*, 11(1):13–34, 2005.

Tara Bedi, Aline Coudouel, and Kenneth Simler, editors. *More than a Pretty Picture: Using Poverty Maps to Design Better Policies and Interventions*. World Bank Publications, September 2007.

Philip Bell. Weighting and Standard Error Estimation for ABS Household Surveys. *ABS Methodology Advisory Committee: Australian Bureau of Statistics*, 2000.

Doorley K. Keane C. Mowlds C. O'Shea R. Bercholz, M. and A. Simon. Distributional Impact of Tax and Welfare Policies: Budget 2026. *QEC Special Article*, 2025. URL `https://www.esri.ie/publications/distributional-impact-of-tax-and-welfare-policies-budget-2026`.

Malcolm Campbell and Dimitris Ballas. Simalba: A Spatial Microsimulation Approach to the Analysis of Health Inequalities. *Frontiers in Public Health*, 4:230, 2016.

Central Statistics Office. Census of Population. `https://www.cso.ie/en/statistics/population/censusofpopulation2022/`, 2022a.

Central Statistics Office. Geographical Profile of Ireland 2022, Household Incomes. `https://data.cso.ie/product/GPII22`, 2022b.

Central Statistics Office (CSO). Number of Families, Persons or Children in Private Households, 2022a. URL `https://data.cso.ie/table/SAP2022T4T1CTY`.

Central Statistics Office (CSO). Estimates of Household Income (RAA02). *Central Statistics Office (CSO), Ireland*, 2022b. URL `https://data.cso.ie/table/RAA02`.

Department for Social Protection. Statistical Information on Social Welfare Services, Annual Report, 2022.

Jean-Claude Deville and Carl-Erik Särndal. Calibration Estimators in Survey Sampling. *Journal of the American Statistical Association*, 87(418):376–382, 1992.

Jean-Claude Deville, Carl-Erik Särndal, and Olivier Sautory. Generalized Raking Procedures in Survey Sampling. *Journal of the American Statistical Association*, 88(423):1013–1020, 1993.

Chris Elbers and Roy van der Weide. Estimation of Normal Mixtures in a Nested Error Model with an Application to Small Area Estimation of Poverty and Inequality. Policy Research Working Paper Series 6962, The World Bank, July 2014. URL `https://ideas.repec.org/p/wbk/wbrwps/6962.html`.

Chris Elbers, Jean O. Lanjouw, and Peter Lanjouw. Micro-Level Estimation of Poverty and Inequality. *Econometrica*, 71(1):355–364, 2003.

Niall Farrell. Small Area Poverty Estimation by Conditional Monte Carlo. *ESRI Working Paper*, WP773, 2024.

Niall Farrell and Stefano Ceolotto. Keeping Our Heads Above Water: Spatially Heterogeneous Social Vulnerabilities and Climate Adaptation. *ESRI Working Paper 771*, 2024.

Niall Farrell, Cathal O'Donoghue, and Karyn Morrissey. Spatial Microsimulation using Quota Sampling. *Teagasc Rural Economy Research Series Working Paper*, 7, 2010.

Niall Farrell, Karyn Morrissey, and Cathal O'Donoghue. Creating a Spatial Microsimulation Model of the Irish Local Economy. In R. Tanton and K. Edwards, editors, *Spatial Microsimulation: A Reference Guide for Users*, volume 6 of *Understanding Population Trends and Processes*. Springer, 2012.

Tim Goedemé, Karel Van den Bosch, Lina Salanauskaite, Gerlinde Verbist, et al. Testing the Statistical Significance of Microsimulation Results: A Plea. *International Journal of Microsimulation*, 6(3): 50–77, 2013.

Tom Goedemé. The Standard Error of Estimates Based on EU-SILC: An Exploration through the Europe 2020 Poverty Indicators. CSB Working Paper Series WP 10/09, Herman Deleeck Centre for Social Policy, University of Antwerp, 2010.

Dilovar Haydarov, Cathal O'Donoghue, Mary Ryan, and Chaosheng Zhang. Local Natural Capital Influences on the Geospatial Distribution of Farm Incomes. *International Journal of Microsimulation*, 17 (1):1–22, 2024.

Jesko Hentschel, Jean Olson Lanjouw, Peter Lanjouw, and Javier Poggi. Combining Census and Survey Data to Trace the Spatial Dimensions of Poverty: A Case Study of Ecuador. *The World Bank Economic Review*, 14(1):147–165, 2000.

Claire Keane, Karina Doorley, Theano Kakoulidou, and Seamus O'Malley. SWITCH: A Tax-Benefit Model for Ireland Linked to Survey and Register Data. *International Journal of Microsimulation*, 16 (1):65–88, 2023.

Paul Kilgarriff, Cathal O'Donoghue, Martin Charlton, and Ronan Foley. Intertemporal Income in Ireland 1996–2011: A Spatial Analysis. *International Journal of Microsimulation*, 9(2):1–23, 2016.

Anders Klevmarken. Statistical Inference in Microsimulation Models: Incorporating External Information. *International Journal of Microsimulation*, 15(1):111–120, 2022.

Robin Lovelace, Mark Birkin, Dimitris Ballas, and Eveline Van Leeuwen. Evaluating the Performance of Iterative Proportional Fitting for Spatial Microsimulation: New Tests for an Established Technique. *Journal of Artificial Societies and Social Simulation*, 18(2), 2015.

Thomas Lumley. *R package `survey`: Analysis of Complex Survey Samples*, 2020. URL `https://CRAN.R-project.org/package=survey`. R package version 4.4-8.

Thérèse McDonnell, Emma Nicholson, Gerard Bury, Claire Collins, Ciara Conlon, Kevin Denny, Michael O'Callaghan, and Eilish McAuliffe. Policy of Free GP Care for Children under 6 Years: The Impact on Daytime and Out-of-hours General Practice. *Social Science Medicine*, 296, 2022.

Nicholas Minot. Generating Disaggregated Poverty Maps: An Application to Vietnam. *World Development*, 28(2):319–331, 2000.

Riyana Miranti, Rebecca Cassells, Yogi Vidyattama, and Justine McNamara. Measuring Small Area Inequality using Spatial Microsimulation: Lessons Learned from Australia. *International Journal of Microsimulation*, 8(2):152–175, 2016.

Angelo Moretti and Adam Whitworth. Estimating the Uncertainty of a Small Area Estimator based on a Microsimulation Approach. *Sociological Methods & Research*, 52(4):1785–1815, 2023.

Karyn Morrissey, Graham Clarke, Dimitris Ballas, Stephen Hynes, and Cathal O'Donoghue. Examining Access to GP Services in Rural Ireland using Microsimulation Analysis. *Area*, 40(3):354–364, 2008.

Karyn Morrissey, Cathal O'Donoghue, Graham Clarke, and Jinjing Li. Using Simulated Data to Examine the Determinants of Acute Hospital Demand at the Small Area Level. *Geographical Analysis*, 45(1): 49–76, 2013.

Ralf Münnich, Rainer Schnell, Hanna Brenzel, Hanna Dieckmann, Sebastian Dräger, Jana Emmenegger, Philip Höcker, Johannes Kopp, Hariolf Merkle, Kristina Neufang, et al. A Population based Regional Dynamic Microsimulation of Germany: the MikroSim Model. *Methods, Data, Analyses: A Journal for Quantitative Methods and Survey Methodology*, 15(2):241–264, 2021.

Anne Nolan and Richard Layte. Growing Up in Ireland: Understanding Use of General Practitioner Services among Children in Ireland. *ESRI Report*, 2017.

Cathal O'Donoghue. *Practical Microsimulation Modelling*. Oxford University Press, 2021.

Cathal O'Donoghue, Dimitris Ballas, Graham Clarke, Stephen Hynes, and Karyn Morrissey, editors. *Spatial Microsimulation for Rural Policy Analysis*. Advances in Spatial Science. Springer, Berlin, Heidelberg, 2013.

Cathal O'Donoghue, Karyn Morrissey, and John Lennon. Spatial Microsimulation Modelling: A Review of Applications and Methodological Choices. *International Journal of Microsimulation*, 7(1):26–75, 2014.

OECD. *Regional Outlook 2023: The Longstanding Geography of Inequalities*. OECD Report, Paris, 2023.

Cathal O'Donoghue, Niall Farrell, Karyn Morrissey, John Lennon, Dimitris Ballas, Graham Clarke, and Stephen Hynes. The SMILE Model: Construction and Calibration. *Spatial Microsimulation for Rural Policy Analysis*, pages 55–86, 2013.

Anastasia Panori, Dimitris Ballas, and Yannis Psycharis. Simathens: A Spatial Microsimulation Approach to the Estimation and Analysis of Small Area Income Distributions and Poverty Rates in the City of Athens, Greece. *Computers, Environment and Urban Systems*, 63:15–25, 2017.

Azizur Rahman, Ann Harding, Robert Tanton, and Shuangzhe Liu. Methodological Issues in Spatial Microsimulation Modelling for Small Area Estimation. *International Journal of Microsimulation*, 3 (2):3–22, 2010.

J. N. K. Rao and C. F. J. Wu. Bootstrap Inference for Sample Surveys. In *Proceedings of the Section on Survey Research Methods*, pages 866–871. American Statistical Association, 1993.

Kenneth Simler. *Pinpointing Poverty in Europe: New Evidence for Policy Making*. World Bank, Washington, DC., 2016.

Avi C. Singh and Chris A. Mohl. Understanding Calibration Estimators in Survey Sampling. *Survey Methodology*, 22:107–116, 1996.

Denisa M. Sologon, Karina Doorley, and Cathal O'Donoghue. Drivers of Income Inequality: What Can We Learn Using Microsimulation? In Klaus F. Zimmermann, editor, *Handbook of Labor, Human Resources and Population Economics*, pages 1–37. Springer International Publishing, 2020.

Holly Sutherland and Francesco Figari. Euromod: the european union tax-benefit microsimulation model. *International journal of microsimulation*, 6(1):4–26, 2013.

Robert Tanton. A Review of Spatial Microsimulation Methods. *International Journal of Microsimulation*, 7(1):4–25, 2014.

Robert Tanton. Spatial Microsimulation: Developments and Potential Future Directions. *International Journal of Microsimulation*, 11(1):143–161, 2018.

Robert Tanton, Paul Williamson, and Ann Harding. Comparing Two Methods of Reweighting a Survey File to Small Area Data. *International Journal of Microsimulation*, 7(1):76–99, 2014.

Tarozzi and Deaton. Using Census and Survey Data to Estimate Poverty and Inequality for Small Areas. *Review of Economics and Statistics*, 9:773–792, 2009.

Laureti Tiziana and Benedetti Ilaria. Measuring and Communicating Uncertainty of Poverty Indicators at Regional Level. *Eurostat Statistical Working Papers*, 2020.

Melanie N Tomintz, Bernhard Kosar, and Victor M García-Barrios. simSALUD: Design and Implementation of an Open-source Wizard based Spatial Microsimulation Framework. *International Journal of Microsimulation*, 10(2):118–143, 2017.

Amaya Vega, Paul Kilgarriff, Cathal O'Donoghue, and Karyn Morrissey. The Spatial Impact of Commuting on Income: A Spatial Microsimulation Approach. *Applied Spatial Analysis and Policy*, 10(4): 475–495, 2017.

Vijay Verma, Achille Lemmi, Gianni Betti, Francesca Gagliardi, and Mario Piacentini. How Precise are Poverty Measures Estimated at the Regional Level? *Regional Science and Urban Economics*, 66: 175–184, 2017.

Yogi Vidyattama and Robert Tanton. Projecting Small Area Statistics with Australian Microsimulation Model (SpatialMSM). *Australasian Journal of Regional Studies, The*, 16(1):99–126, 2010.

Quoc Ngu Vu and Robert Tanton. The Distributional and Regional Impact of the Australian Government's Household Stimulus Package. *The Australasian Journal of Regional Studies*, 16(1):127–145, 2010.

Adam Whitworth, Eleanor Carter, Dimitris Ballas, and Graham Moon. Estimating Uncertainty in Spatial Microsimulation Approaches to Small Area Estimation: A New Approach to Solving an Old Problem. *Computers, Environment and Urban Systems*, 63:50–57, 2017.

Alex Whitworth. `synthACS`: Spatial Microsimulation Modeling with Synthetic American Community Survey Data. *Journal of Statistical Software*, 104:1–30, 2022.

Jesse Wiki, Lukas Marek, Chris Sibley, and Daniel Exeter. Estimating Quality of Life: A Spatial Microsimulation Model of Well-being in Aotearoa New Zealand. *Social Science & Medicine*, 330: 116054, 2023.

Kirk M. Wolter. The Jackknife Method. In *Introduction to Variance Estimation*, Statistics for Social and Behavioral Sciences, chapter 4, pages 151–193. Springer, New York, NY, 2 edition, 2007.

Guoqiang Wu, Alison Heppenstall, Petra Meier, Robin Purshouse, and Nik Lomax. A Synthetic Population Dataset for Estimating Small Area Health and Socio-Economic Outcomes in Great Britain. *Scientific Data*, 9(1):19, 2022.