How does attrition affect estimates of persistent poverty rates? The case of European Union statistics on income and living conditions (EU-SILC)

S.P. JENKINS AND P. VAN KERM

2017 edition

STATISTICAL
WORKING PAPERS

eurostat

How does attrition affect estimates of persistent poverty rates? The case of European Union statistics on income and living conditions (EU-SILC)

S.P. JENKINS AND P. VAN KERM

2017 edition

Luxembourg: Publications Office of the European Union, 2017

ISBN 978-92-79-64143-5
ISSN 2315-0807
doi: 10.2785/86980
Cat. No: KS-TC-16-025-EN-N

Theme: Population and social conditions
Collection: Statistical working papers

Preface

Eurostat is the Statistical Office of the European Union (EU). Its mission is to provide high-quality statistics on Europe. To that end, it gathers and analyses data from the National Statistical Institutes (NSIs) across Europe and provides comparable and harmonised data for the EU to use in the definition, implementation and analysis of EU policies. Its statistical products and services are also of great value to Europe's business community, professional organisations, academics, librarians, NGOs, the media and citizens.

In the field of income, poverty, social exclusion and living conditions, the EU Statistics on Income and Living Conditions (EU-SILC) is the main source for statistical data at European level.

Over the last years, important progress has been achieved in EU-SILC as a result of the coordinated work of Eurostat and NSIs.

In June 2010, the European Council adopted a social inclusion target as part of the Europe 2020 Strategy: to lift at least 20 million people in the EU from the risk of poverty and exclusion by 2020. To monitor progress towards this target, the 'Employment, Social Policy, Health and Consumer Affairs' (EPSCO) EU Council of Ministers agreed on an 'at risk of poverty or social exclusion' indicator. To reflect the multidimensional nature of poverty and social exclusion, this indicator consists of three sub-indicators: i) at-risk-of-poverty (i.e. low income); ii) severe material deprivation; and iii) (quasi-)joblessness.

In this context, the Second Network for the Analysis of EU-SILC (Net-SILC2) is bringing together NSIs and academic expertise at international level in order to carry out in-depth methodological work and socio-economic analysis, to develop common production tools for the whole European Statistical System (ESS) as well as to ensure the overall scientific organisation of the third and fourth EU-SILC conferences.

It should be stressed that this methodological paper does not in any way represent the views of Eurostat, the European Commission or the European Union. This is independent research which the authors have contributed in a strictly personal capacity and not as representatives of any Government or official body. Thus they have been free to express their own views and to take full responsibility both for the judgments made about past and current policy and for the recommendations for future policy.

This document is part of Eurostat's Methodologies and working papers collection, which are technical publications for statistical experts working in a particular field. These publications are downloadable free of charge in PDF format from the Eurostat website: http://ec.europa.eu/eurostat/en/web/products-statistical-working-papers.

Eurostat databases are also available at this address, as are tables with the most frequently used and requested short- and long-term indicators.

Abstract[1]

Among the primary indicators of social inclusion is the persistent at risk of poverty rate, defined as the proportion of persons in a country who are at risk of income poverty in the current year and who were at risk of income poverty in at least two of the preceding three years. Evidence about poverty persistence is an important complement to information about poverty prevalence at a point in time. Estimates of persistent at risk of poverty rates are derived from the longitudinal component of EU SILC in which the fortunes of individuals are tracked over four consecutive years, in principle. In practice, not all of the individuals present in the first sample year provide four years of income data: there is attrition and estimates of persistent at risk of poverty measure may therefore not be reliable. Rates of attrition from the four-year EU-SILC samples used to calculate persistent poverty rates vary substantially across Member States, and there is also substantial cross-national diversity in the characteristics of individuals lost to follow-up. This working paper documents such patterns in detail and provides evidence that application of longitudinal weights does not fully account for the effects of attrition, and that different assumptions about the poverty status of attritors lead to wide bounds for estimates of persistent poverty rates for most Member States.

Acknowledgements

This work has been supported by the second Network for the analysis of EU-SILC (Net-SILC2), funded by Eurostat, and partially supported by core funding of the Research Centre on Micro-Social Change at the Institute for Social and Economic Research by the University of Essex and the UK Economic and Social Research Council (award ES/L0009153). The European Commission bears no responsibility for the analyses and conclusions, which are solely those of the authors. We thank Tony Atkinson, Carlos Farinha Rodrigues, Anne-Catherine Guio, Eric Marlier, Veli-Matti Törmälehto, and participants at the Net-SILC2 Lisbon conference (October 2014) for comments and suggestions.

(1) Authors addresses: Stephen P. Jenkins, Department of Social Policy, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, U.K. Email: s.jenkins@lse.ac.uk
Philippe Van Kerm, Luxembourg Institute of Socio-Economic Research, 11, Porte des Sciences, L-4366 Esch-sur-Alzette, Luxembourg. Email: philippe.vankerm@liser.lu

Table of contents

1. Introduction

Over the last decade and through its Open Method of Coordination, the European Union (EU) has agreed a set of common objectives for monitoring and measurement of social protection and social inclusion, together with a set of indicators to assess national and EU progress towards these goals. Among the primary indicators of social inclusion is the persistent at-risk-of-poverty rate, defined as the proportion of persons in a country who are at risk of income poverty in the current year and who were at risk of income poverty in at least two of the preceding three years. Evidence about poverty persistence is an important complement to information about poverty prevalence at a point in time: it is widely agreed that poverty is worse for an individual, the longer he or she experiences it. Eurostat derives estimates of persistent at-risk-of-poverty rates using samples from the longitudinal component of the EU Statistics on Income and Living Conditions (EU-SILC) in which the fortunes of individuals are tracked over four consecutive years. Because not all of the individuals present in the first sample year provide four years of income data – there is attrition – estimates of persistent at-risk-of-poverty measure may not be reliable. In this paper, we analyse the extent to which this is the case, and how the potential problems vary across EU member states.

Attrition is a potential problem for two reasons. First, it means that the sample size for the four-year sample used to calculate a persistent at-risk-of-poverty rate is smaller than the size of the sample of respondents in the first year of the four (wave 1), and a smaller sample size leads to *less precise inference* (larger standard errors and wider confidence intervals). Second, if the individuals who are lost to follow-up differ systematically from the initial respondent sample – the case of non-random or 'differential' attrition – the four-year sample may not be representative of the underlying population, thereby leading to *biased estimates* of persistent at-risk-of-poverty rates.

The longitudinal weights supplied with EU-SILC longitudinal data are intended to address the second problem. The idea is that, if differences in the chances of sample dropout can be fully characterised in terms of differences in individuals' observed characteristics, weighting will make the four-year sample representative of the initial sample. Individuals with characteristics associated with large dropout probabilities receive relatively large weights to compensate for the large fraction of similar individuals that have been lost. Individuals less likely to dropout receive relatively small weights. The weighting strategy works as long as observable characteristics predict dropout probabilities well and those who remain in the sample are not systematically different from those who attrit. However, problems arise if the chances of attrition also depend on unobserved characteristics that are systematically correlated with the chances of being persistently at-risk-of-poverty. Because such characteristics are unobserved, their impact is difficult to assess.

We provide indirect evidence about attrition bias, including application of a novel method that places bounds on estimates of persistent at-risk-of-poverty rates. Application of a parametric model incorporating strong assumptions about the relationship between attrition and persistent at-risk-of-poverty (a 'bivariate probit model with selection') turns out to be uninformative.

Our research builds on analysis of attrition in EU-SILC's forerunner, the European Community Household Panel (ECHP), undertaken by Behr, Bellgardt, and Rendtl (2005) and Watson (2003). EU-SILC differs substantially from the ECHP which ran between 1994 and 2001. Although both sources employ annual data collection, EU-SILC longitudinal data refer to information collected over a four-year period, rather than up to eight years. Instead of using a survey instrument with a cross-nationally harmonised design (household panel surveys in ECHP), EU-SILC uses output harmonisation. Countries are mandated to deliver a number of statistics conforming to particular specifications (and the data used to create them) but have some discretion about the ways in which the information is collected. Most notably, some countries use household panel surveys to collect the longitudinal data; others use linked administrative registers.[2] In addition, there are many more countries contributing EU-SILC data than were in the ECHP: we use 23 countries in our analysis; there were only 15 countries covered by the ECHP.

Behr, Bellgardt, and Rendtl (2005) and Watson (2003) both drew attention to a substantial diversity in response rates in ECHP and, moreover, their conclusions were that, although the amount of attrition was relatively large, its effects on estimates of poverty rates and quintile transition probabilities were relatively benign. Indeed, Watson went so far as to state that 'fears that attrition has undermined the representativeness of the ECHP samples in later waves of the survey are largely unfounded' (2003: 361). Her results about representativeness are similar to those reported by Fitzgerald, Gottschalk, and Moffitt (1998) for the US Panel Study of Income Dynamics.

Patterns of attrition and their consequences may have changed substantially over the last decade. Also, with many more countries with data, and output harmonisation rather than input harmonisation, there is much greater scope for differences across Member States. Our analysis of attrition and estimation of persistent at-risk-of poverty rates in EU-SILC data is therefore not only timely but also important given the place of this indicator in the EU's portfolio of social inclusion indicators.

(2) On this, see e.g. Lohmann (2011) or, for a thorough discussion, Jäntti, Tömälehto, and Marlier (2013).

The remainder of the paper is organised as follows. In Section 2, we explain the data that we use, drawn from the 2011 longitudinal EU-SILC User DataBase. This discussion covers the definition of the persistent at-risk-of-poverty rate, how attrition arises, the weights that are available, and the samples that we use in the analysis. The extent of attrition across Member States, and how it varies with personal characteristics, is described in Section 3. In Section 4, we analyse the implications of sample dropout, again contrasting the situation across countries. We assess effects on representativeness by comparing estimates of at-risk-of-poverty rates from the full initial sample with estimates derived from the smaller four-wave sample. We look more directly at the impact of attrition, first by using a parametric model and, second, by estimating bounds for the persistent at-risk-of-poverty measure making no assumptions about the relationship between attrition and poverty. Section 5 contains a summary and conclusions.

2. Data, definitions, sample selection, weighting

Our analysis is based on the 2011 EU-SILC longitudinal files. More specifically we use the scientific-use release of the longitudinal EU-SILC files made available to the NetSILC-2 project, which are an update of UDB 2011-1, released August 2013. These files refer to data covering the four survey years 2008–2011. Because the reference period for EU-SILC income data is the calendar year preceding the year of data collection, the income years covered are 2007–2010.[3]

2.1 At-risk of-poverty rates and persistent at-risk-of-poverty rates

Following EU official definitions, an individual's 'at-risk-of-poverty' status in a given income year is determined by the equivalised household disposable income of the household to which he or she belongs. (For further details of the sources included in household income and the equivalence scale, see Eurostat (2010).) A person is counted as being at-risk-of-poverty (henceforth poor) in a given year if his or her equivalised household disposable income is less than 60 per cent of the national median equivalised household income for that year.[4] The current at-risk-of-poverty rate (henceforth *current poverty rate*) for a particular country or group within a country is the proportion of persons in that country or group who are poor in the current income year.

The persistent at-risk-of-poverty rate (henceforth *persistent poverty rate*) is the proportion of persons in the country or group who are currently poor and who were poor in at least two of the preceding three years. Thus in our longitudinal data, the persistent poverty rate refers to the proportion of individuals who were poor in 2010 as well as in at least two of the three previous years (2007–2009). This indicator is the principal official EU indicator on social inclusion for which estimation is based on the longitudinal component of EU-SILC, and hence the indicator that is most sensitive to attrition issues.

2.2 Samples

EU-SILC has a four-year rotating panel design. A fresh sample of households is drawn every year in every country, and respondents in this sample are eligible for interview in each of the following three years, contributing a total of up to four interviews. In each calendar year, data are available from four cohorts of respondents and contribute to the EU-SILC cross-section data. The 2011 EU-SILC longitudinal data (and similarly in preceding releases) consist of the three subsamples that provide data in 2011 and in at least one earlier survey year as well, i.e. the cohorts that entered the survey in 2008, 2009, or 2010.

To examine the magnitude and pattern of attrition, and to assess their implications for estimation of persistent poverty rates, we work with the 2008 rotation group sample which provide data over up to four years and is therefore the basis for calculation of the official 2011 persistent at-risk-of-poverty indicator. We use the samples for 23 countries: we exclude the samples for Luxembourg (because no rotation group was started in 2008), Norway (because some of the relevant sample weights were not available – see below), Denmark (because the 2011 database appears to exclude households that attritted before the fourth interview), and Sweden (because of unexplained differences in sizes between the 2010 and 2011 versions of the 2008 rotation group samples).

Our examination of the magnitude and effect of attrition relies on two overlapping subsamples. The first sample is composed of all individuals from all households in the rotation groups for survey years 2008–2011 that responded at wave 1 (wave 1 is the year in which households entered the survey, i.e. 2008) irrespective of their subsequent participation. For a number of countries, this corresponds to all households recorded in the household registry (the 'd-file'); for other countries, we discard the households from the household registry that are recorded as not participating even at wave 1. We refer to this full sample of the 2008 rotation group as the full *W1 Sample*. In principle, this sample should provide estimates close to those derived from the full 2008 cross-sectional sample. We return to this point later.

Our second subsample is composed of the subset of individuals from the W1 Sample that belong to a household succesfully interviewed in each one of the four survey years 2008–2011. This is the *four-wave Balanced Sample*, from which persistent poverty rates can be calculated.[5] We consider only individuals who were living in a household that was interviewed at wave 1: we discard children born after wave 1 as well as co-residents that joined a sample household after wave 1 since, by construction,

(3) Data for the United Kingdom deviate from this rule: the income reference period refers to the period around the date of interview, and income totals are converted subsequently to annual equivalents pro rata. Ireland also deviates from the rule with income data referring to the 12 month period prior to the interview. Data for Ireland were not released in the 2011 longitudinal EU-SILC.

(4) Throughout our analysis, poverty lines for each country and year are taken from Eurostat (2014). These thresholds are derived from the cross-sectional EU-SILC datasets. Because of the rotating panel structure of the EU-SILC data (see below), the cross-section samples are much larger than any longitudinal sample. Poverty lines provided by Eurostat (2014) are therefore more accurately estimated than those that could be computed from our longitudinal samples.

(5) More precisely, it is based on the subsample with valid (non-missing) data on household income in the EU-SILC data files in all years. However, because missing information on income is imputed (and we use the imputed values), all households contain non-missing data on income in the EU-SILC data files.

these individuals do not have a full four-year set of responses. An important distinction between 'register' and 'survey' countries then comes into play. Two distinct models are used in EU-SILC in the follow-up of respondents. Survey countries use a standard longitudinal survey design and aim to follow over time all household members initially interviewed; that is, if an original household splits, they attempt to follow all individuals in all the newly-formed households. By contrast, register countries use a 'selected respondent' design and only track ***one*** 'selected respondent' from each original household. Only co-residents who remain living in the same household as the selected respondent provide information on income over time. This following rule mechanically leads to higher attrition rates in register countries, as we show below, since not all of the households that split are tracked.

Variations in practice and in he success of tracking of individuals and interviewing 'split-off' households has been shown to vary widely across countries by Iacovou and Lynn (2013). These are likely sources of the cross-country differences in attrition rates documented below.

2.3 Attrition

The differences in size and composition between the W1 Sample and the four-wave Balanced Sample reflect attrition. Not all individuals or households eligible for an interview after the first interview provide data in subsequent years. There are four reasons for this.

The first is related to the following rule used by register countries: co-residents of the main 'selected respondent' that leave a household are not followed, by design. Second, some individuals or households move out of scope after the first interview, for example because they die, or move abroad permanently, or move into an institution. Third, eligible individuals may not be followed by the data collection agency, or the agency may be unsuccessful in tracking them (with the chances greater for individuals that split off from a household, or where all members of a household move from the initially-sampled address). Fourth, individuals or households may refuse to participate in the survey in the second interview or later.[6]

The first kind of attrition among register countries – attrition by design – is not necessarily a problem per se; the issue in the current context is that it leads to cross-national inconsistencies in EU-SILC. The second kind of sample dropout reflects the dynamics of a population and is a natural feature that is built into the data collection design (based on representation of the population of individuals in private households in a particular country). By contrast, the third and fourth types of attrition are undesirable and, other things being equal, data collection agencies should aim to minimise them. Country-specific factors may also play a role, for example, whether up-to-date address registers exist, the prevalence of geographical mobility by households, general attitudes towards surveys, etc.

2.4 Sampling weights

Sampling weights are designed to adjust for biases arising from cross-sectional non-response and subsequent longitudinal attrition. The EU-SILC longitudinal files include five types of sample weights (Museux 2006), of which two are relevant to our analysis.[7]

The first set of weights is the *individual-level base weights* (variable *rb060*). In wave 1, this is the design weight adjusted for non-response and calibrated. In later waves, it is the base weight of previous year adjusted for non-response. When individuals leave the sample, they are attributed with a weight of zero for each wave thereafter. Our analysis of the W1 Sample uses *rb060* to ensure the sample accounts for non-proportional sampling design (and initial non-response), and for differential attrition, and is calibrated to population totals in 2008.

The second set of weights that we use, *rb064*, is the individual-level longitudinal weights created for analysis of data for the four survey years 2008–2011 and, of course, the weights are only relevant for the single rotation group that provides data for these four years. For analysis of the four-wave Balanced Sample, we contrast results obtained with *rb064* (constructed to ensure that the balanced sample remains representative of the original 2008 population), with *rb060* at their 2008 value (so they correct for initial non-response and sampling rates, but not for differential attrition), and with *rb060* at their 2011 values (in which case they are similar to *rb064*).

[6] Iacovou and Lynn (2013) discuss the difficulty in consistently identifying the causes of attrition in EU-SILC across the different countries.
[7] We do not use household-level cross-section weights (db090), individual-level longitudinal weights applicable only to analysis involving data for 2010 and 2011 for the three rotation groups providing data in those two years (rb062), nor the individual-level longitudinal weights applicable only to analysis of data for 2009-2011 for the two rotation groups providing data for those two years (rb063).

We also create our own bespoke set of longitudinal weights (discussed later). The advantage of these weights is that we can use them to engage in a number of counterfactual exercises that we cannot undertake with the weights that are supplied. We show below that these weights generally closely reproduce estimates derived using the official longitudinal weights although our bespoke weights are derived using variables available in the longitudinal data files, and we do not have access to all the factors employed by statistical offices when producing longitudinal weights (*rb064*), nor do we attempt to calibrate our weights to known population totals, for example, as derived from other data sources or from the full EU-SILC cross-section files.

3. How much attrition is there? Who drops out?

In this section, we document how much attrition there was in the 2008–2011 EU-SILC longitudinal data, and which types of individual were most likely to be lost to follow-up. We discuss attrition – or its complement, sample retention – in terms of differences between the full Wave 1 Sample and the smaller four-wave Balanced Sample.

3.1 How much attrition is there overall?

The overall retention rate for each country is the fraction of the country's full W1 sample that belongs to the Balanced Sample. More precisely we calculate the retention rate as the proportion of individuals belonging to a respondent household at wave 1 which remains in a respondent household in each of the three subsequent waves. These rates are reported in Figure 1, panel (a), with the absolute numbers of individuals in each country's samples shown in panel (b).

There are very large differences in retention rates across countries, ranging from greater than 90 per cent to nearer 40 per cent. The UK stands out as having a particularly low retention rate, nearly 10 percentage points smaller than the next smallest rate, 50 per cent for Slovenia. There is a cluster of three countries with remarkably large retention rates: those for Romania and Bulgaria are all near 90 per cent. Unsurprisingly, the method of data collection is related to the retention rate of original household members: 'register' countries (identified by the circles in Figure 22.1; Slovenia, Finland, Iceland and the Netherlands) tend to exhibit comparatively low retention rates for reasons outlined above.

Figure 1, panel (b), shows that there are substantial differences across countries in the numbers of individuals in the Wave 1 Samples. Three countries have samples of more than 10 000 individuals (Slovenia, Spain, Poland, and Italy), and four countries have Wave 1 Samples of fewer than 3 000 individuals (Cyprus, Malta, and Iceland). The numbers of individuals in the four-wave Balanced Samples are smaller of course. The maximum sample size is around 9 000 (Italy) and 12 of the 23 countries have samples with fewer than 3 000 individuals.

Figure 1: Retention rates and sample sizes by country, 2008-2011

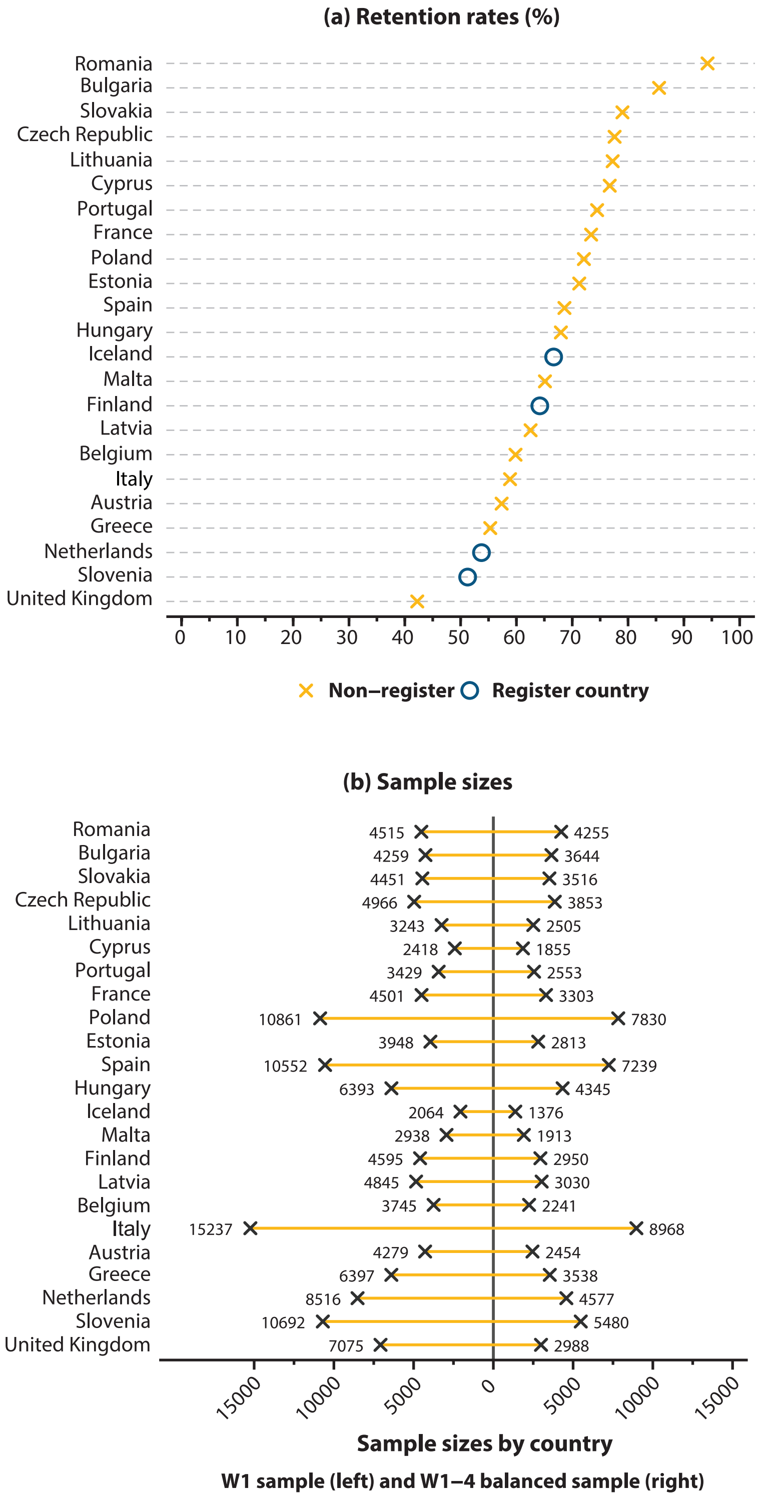

Note: The retention rate is the proportion of individuals belonging to a respondent household at Wave 1 (2008) which remains in a participating household in each of the three subsequent waves. Only these individuals are used for the calculation of the 2011 persistent poverty rates. Unweighted proportions of wave 1 sample.

Source: Calculations from 2011 EU-SILC Longitudinal data; 2008 rotation group only.

3.2 Attrition's effect on the precision of estimates

The decrease in sample sizes associated with attrition means that, questions of representativeness and hence bias aside (on which see below), estimates of persistent poverty rates are estimated less precisely. Standard errors are larger, and confidence intervals are wider. The effects of differences in sample size on the sampling variability of estimates can be gauged by noting that the persistent poverty rate is a proportion (p), and there is a standard formula for the standard error of a proportion. The standard error of p is given by $d\sqrt{p(1-p)/N}$, where N is the sample size and d is a design effect arising because of the complex survey design. Figure 2 plots standard errors as a function of N in the range observed in EU-SILC longitudinal data for values of p which cover the range of estimates observed for the persistent poverty rates (0.05, 0.10, 0.15, and 0.20). On the basis of estimates reported for the persistent poverty rates by Goedemé (2013), we set d equal to 1.8, i.e. survey design effects such as stratification and clustering (e.g. of individuals into households, and households into primary sampling units) increase the standard error by 80 per cent compared to the standard error for a simple random sample of the same size.[8]

On the one hand, Figure 2 may provide some cheering news for analysts. Even with substantial attrition and hence relatively small sample sizes, standard errors for persistent poverty rates at the national level may be sufficiently precise. For example, if the persistent poverty rate is around 20 per cent (see the dotted line) and the sample size is 2 500, the standard error for the rate is around 0.015, so the estimated rate is more than ten times larger than its standard error, and the 95% confidence interval is roughly [17%, 23%]. If the sample size were instead 1 000, then the standard error increases to around 0.025, so the ratio of estimate to standard error is around 8, and the confidence interval is approximately [15%, 25%]. If, instead, the persistent poverty rate is only 5 per cent, then a sample size of 1 000 implies a standard error of around 0.012, so the ratio of estimate to standard error falls to just over 4. Ratios of around 2 or more are often interpreted as indicating statistical significance of sufficient degree.

Figure 2: How standard errors for persistent poverty rates vary with sample size
(standard error of proportion p)

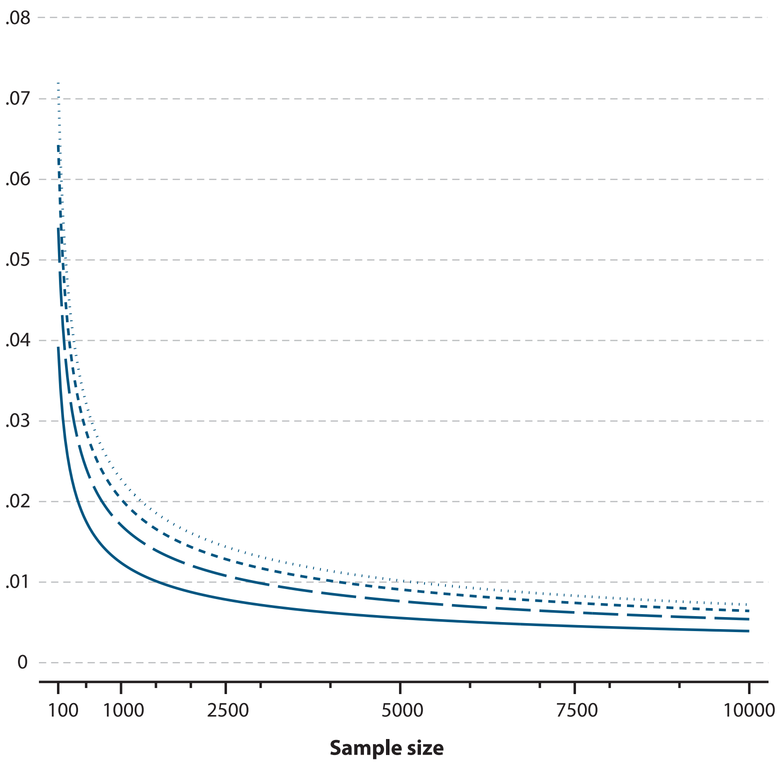

Note: The persistence poverty rate is a proportion.The standard error of a proportion p, SE(p, N) = $d\sqrt{p(1-p)/N}$, where N is the sample size and d is the design effect arising because of the complex survey design. The figure shows SE(p, N) as a function of N in the range observed in EU-SILC longitudinal data for values of p = 0.05, 0.10, 0.15, and 0.20 which cover the range of estimates observed for the persistent poverty rates. We set d = 1.8 (see main text).

Source: Calculations from 2011 EU-SILC Longitudinal data; 2008 rotation group only.

(8) For the sake of argument, we assume here that d is constant across countries. We refer later to situations in which it may not be.

On the other hand, Figure 2 also provides a warning to analysts that estimates of persistent poverty rates for subgroups within a country may not be precisely estimated. For subgroups of particular policy interest, for example individuals living in households headed by a lone parent, sample sizes are likely to number a few hundred at most. With a sample size of 100 and a persistent poverty rate of 20 per cent, the standard error is around 0.06, implying a ratio of estimate to standard error of just over 3 and a 95% confidence interval of approximately [8%, 32%] which is rather wide. In this case, it would be hard to detect statistically significant changes over time in the subgroup persistent poverty rate. The same problem would arise if the persistent poverty rate were smaller than 20 per cent. To add to this cautionary note, we should say that we suggest later (section 4) that, even for large sample sizes (such as for countries as a whole), confidence intervals may be sufficiently wide to encompass differences between estimates that are unbiased and those that are biased because of differential attrition.

3.3 Who drops out? Univariate analysis

We now consider which types of individuals are most likely to be included in the four-wave Balanced Samples. First, we classify individuals according to their characteristics when measured in Wave 1, and calculate retention rates separately for subgroups defined by those characteristics. The individual characteristics we use are poverty status, quintile group of equivalised disposable household income, age and sex, household type, labour market activity status and education level of the household head, and whether the interview questionnaire was completed by a proxy respondent (another household member filling out the questionnaire on behalf of the target respondent). Second, we use probit regression analysis to examine the association between the probability of retention and each characteristic. Subgroups defined by the characteristics mentioned overlap. For example, an individual of pension age is likely to be living in a household headed by someone over the age of 60, and to be retired. Multivariate analysis helps tease out the associations between retention rates and a particular characteristic, holding other characteristics constant. Differences in attrition (retention) rates associated with individual characteristics exemplify the process of *differential* attrition (retention).

Figure 3 shows the univariate breakdowns for each country. Each panel of the figure has a common format. Each overall national retention rate is shown as a cross and subgroup retention rates are shown separately using a numerical code for each subgroup. If subgroup rate for a country is close to the national rate, then attrition is not associated with subgroup membership. Countries are ordered vertically in each chart by their overall retention rate, as in Figure 1.

For example, in Figure 3 panel (a), individuals classified as poor at Wave 1 are coded '11' (the code for non-poor is '10'). It can be seen that poor individuals are more likely to be lost to follow-up in around half the countries and, in a few countries, the differences from the national average are very large. For example, in Belgium and Iceland, poor individuals have a retention rate more than 10 percentage points less than the overall national retention. The difference is about 6 percentage points in the UK. Panel (b) tells a similar story. Retention rates do not vary greatly with income group, except that in a small number of countries, individuals in the poorest fifth are more likely to be lost. (The effects are more muted than in panel (a), probably because the poorest fifth includes more people than are counted as poor.)

Figure 3, panel (c), shows that, in the vast majority of countries, young men (aged between 18 and 40 years) are more likely to attrit than the national average rate, as well as (to a lesser extent) young women. The differences in retention rates across age-sex groups is particularly marked in some countries. For example, in Malta and the UK, the range is around 20 percentage points between the smallest and largest rates.

Figure 3, panel (d), shows that, for many but not all countries, there are large differences in retention rates between household types in some countries, some of which are larger than shown in panel (c). The general picture is that single adult households (with and without) children are most likely to be lost to follow-up, whereas single or couple households with the head aged 60+ have substantially higher retention rates. These differentials are what one would expect given the positive correlation between geographical mobility and age. But dispersion in retention rates by household type is not inevitable: observe the relatively small differentials for the countries with large overall retention rates (at the top of the figure).

Figure 3, panel (e), shows that for most countries retention rates do not vary substantially with the labour market activity status of the household head, though there is a tendency for individuals with unemployed household heads to be more likely to be lost to follow-up and individuals with a retired household head to be less likely to be lost. (In both cases, the head may be the individual him- or herself.) This pattern is particularly marked in some countries. For example, in the UK, the retention rate is just below 20 per cent for individuals with an unemployed head but around 50 per cent for individuals with a retired head (a difference of some 30 percentage points). The corresponding differential is more than 20 percentage points in Malta.

Figure 3: Retention rates, by characteristics
(%)

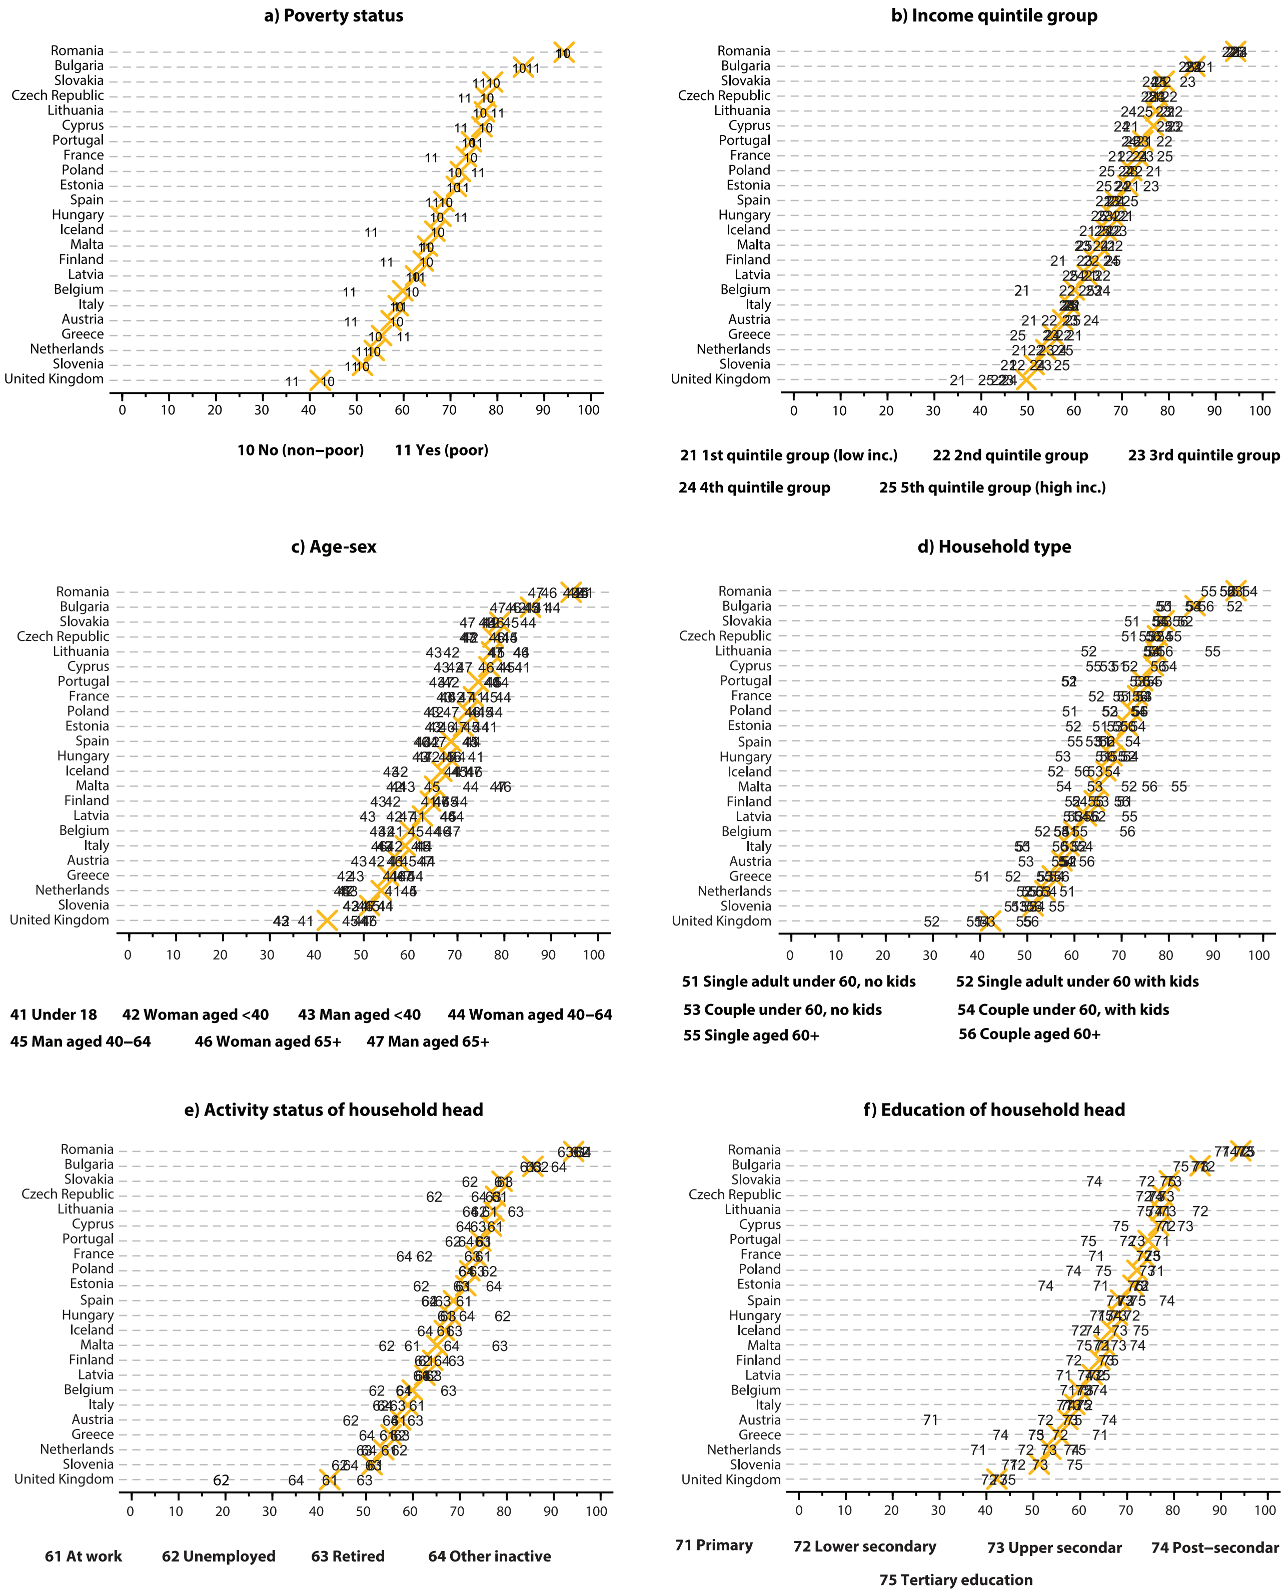

g) Proxy interview

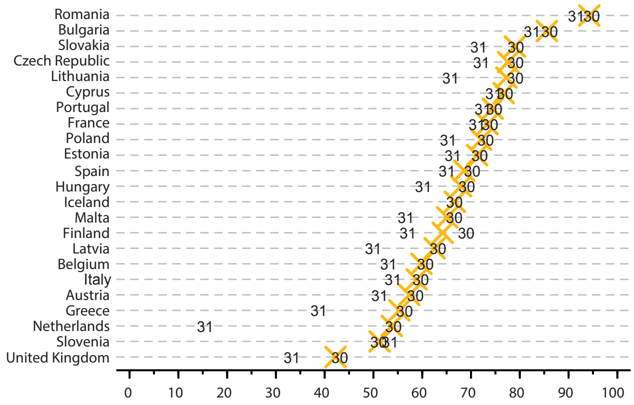

Note: Breakdowns are based on data observed in wave 1. Unweighted proportions of wave 1 sample. Crosses indicate the overall retention rate while numbers identify subgroup retention rates. Retention rates are defined as in Figure 1.

Source: Calculations from 2011 EU-SILC Longitudinal data, 2008 rotation group only.

There appears to be a more complex association between the education level of the household head and retention rates. For countries with relatively low overall retention rates, it is individuals whose household head has either of the two lowest educational levels who have the largest attrition rate. (Austria stands out as an example of this.) And for countries with relatively high overall retention rates, it is individuals whose household head has either of the two highest educational qualifications who have the largest attrition rate. (Look at the cases of Estonia or Slovakia, for example.)

Besides individual or household characteristics, fieldwork-related features are correlated with attrition. Figure 3, panel (g), shows that the retention rate for individuals for whom data were collected from a proxy respondent in wave 1 tend to have low retention rates. This is particularly strong in the Netherlands or Greece, for example. This is unsurprising because a proxy interview in the first wave is indicative of difficulties in securing a respondent's participation to start with. Under-representation of the 'proxy respondent' characteristic itself is unlikely to be a concern; rather, the concern is the extent to which being a 'proxy respondent' is associated with other relevant individual characteristics.

3.4 Who drops out? Multivariate analysis

We now turn to the multivariate analysis of the correlates of retention propensities. We fit probit regressions for the individual probability of retention in the four-wave Balanced Sample, separately by country, using as predictors the variables characterising subgroups that were employed in the univariate analysis albeit with one modification.[9] We do not use information about poverty status because it overlaps closely with income group membership. Regressions are weighted by the 2008 base weight *rb060*.

Regression coefficient estimates are reported in Figure 4, using separate panels to display the results for the various sets of predictor variables. Estimates that do not differ statistically from zero are shown in light grey, and the value of zero (corresponding to no association) is also shown for reference. Values of a coefficient estimate that are less than zero indicate that the relevant predictor is associated with a lower retention probability, and vice versa for coefficients that are greater than zero. When comparing estimates across countries, we focus on the sign and statistical significance of coefficients (as did Behr, Bellgardt, and Rendtl 2005). We shy away from comparing the *magnitude* of a particular regression coefficients across countries because of a well-known problem in the case of binary regression model with comparing coefficients across samples (discussed in section 3.5 and footnote 9 below). Note also that for categorical explanatory variables (e.g. household type), the coefficient for a particular category shown summarises the effect of being in a particular category rather than the reference category (e.g. being single aged 60+ rather than a member of the 'under 18' group). The coefficient for the reference category is zero by construction.

(9) Behr, Bellgardt, and Rendtl (2005) fitted binary logit regression models for retention, applied to each ECHP wave separately, though they also pooled countries and waves in some analyses. Watson (2003) fitted a discrete-time logistic model of the number of waves until attrition from the base sample, but pooled the data from all countries (country differences correspond to intercept shifts).

Figure 4: Coefficients on predictors of retention probabilities

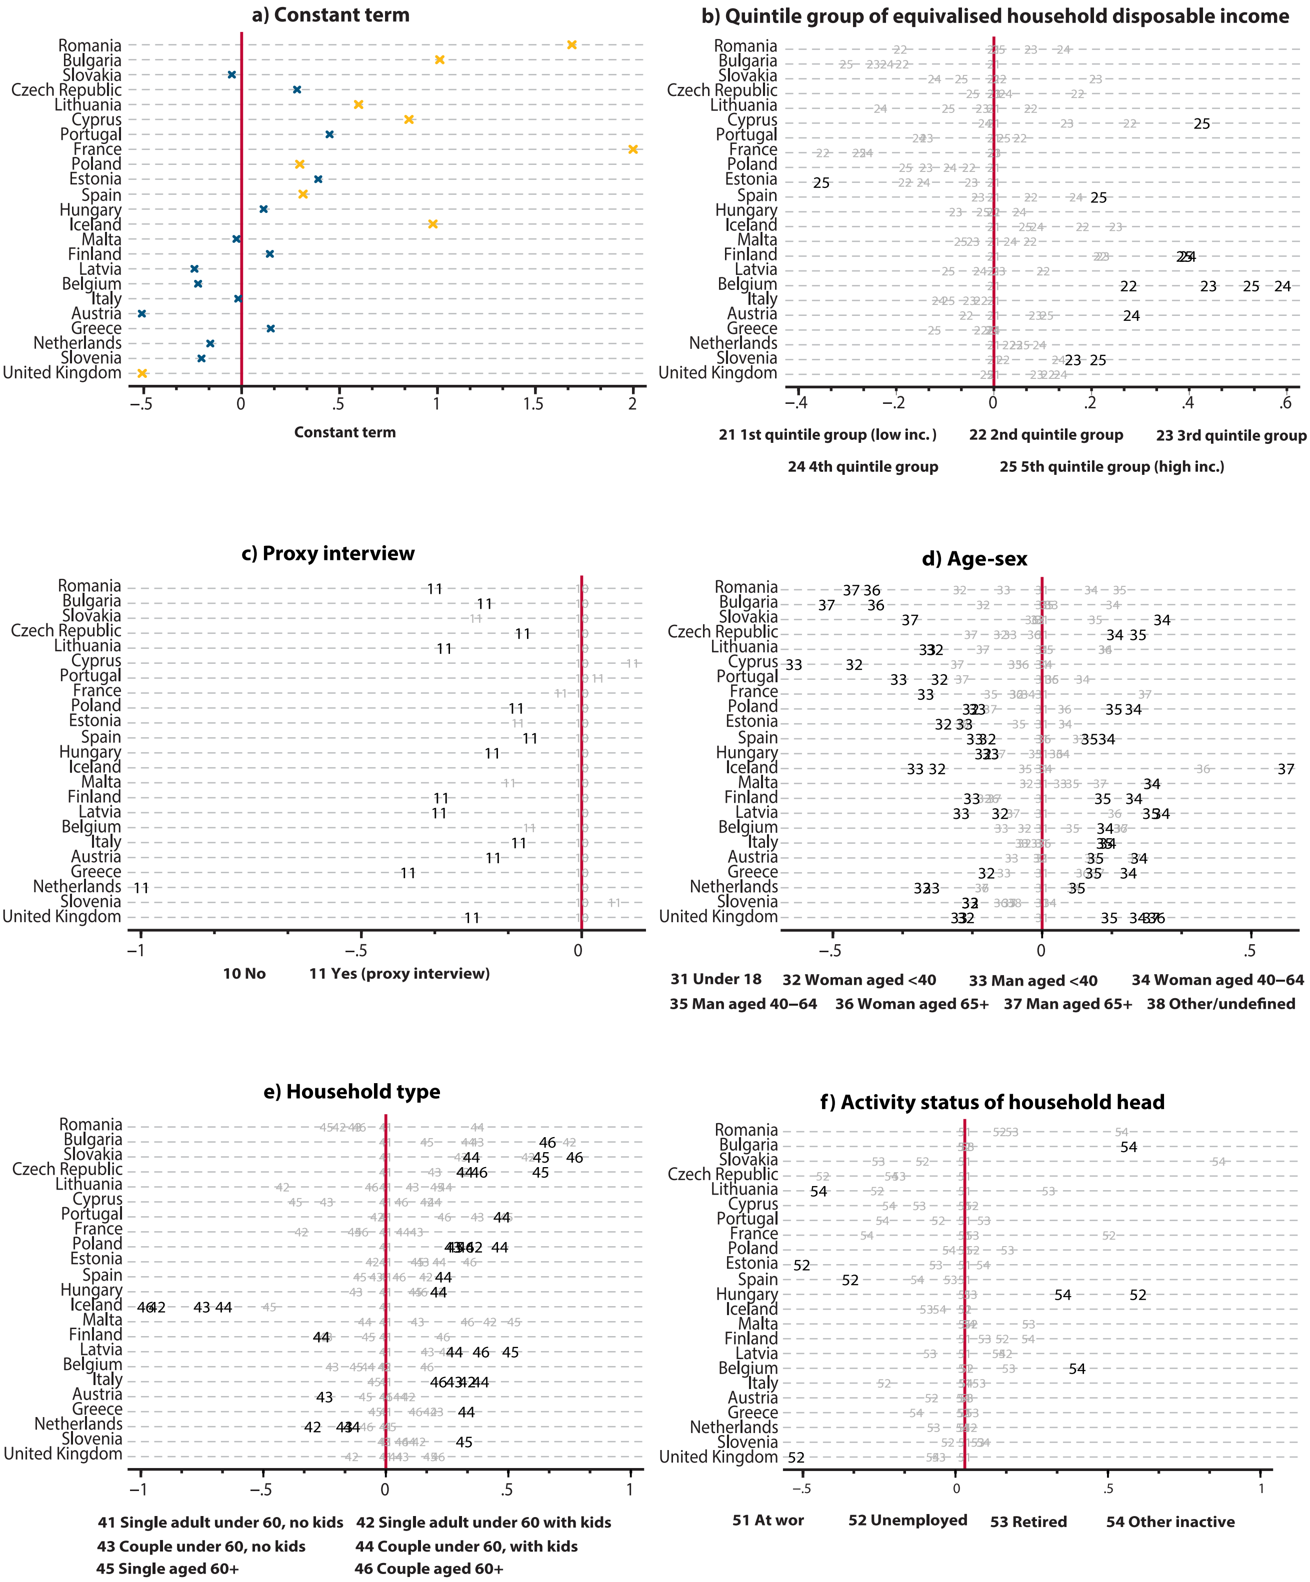

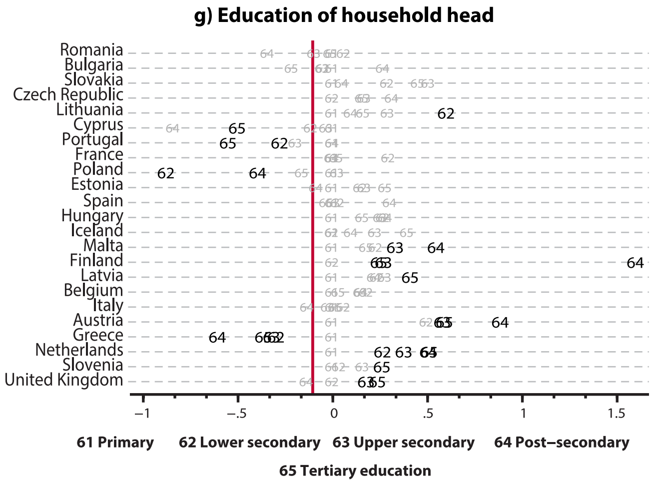

Note: Covariates are based on data observed at wave 1. Retention regressions are weighted by Wave 1 base weights and account for clustering at the household level (but not other sampling design features). Coefficients not significantly different from zero are shown in light grey. Negative coefficients indicate lower retention probability than omitted category. Proxy interview is 0 for children under 16 and non-selected household members in register countries.

Source: Calculations from 2011 EU-SILC Longitudinal data, 2008 rotation group only.

Many of the associations uncovered by the univariate analysis are also found in the multivariate analysis. In every country, individuals with proxy interviews have lower retention probabilities (other things being equal), with the differential between this group and those with full interviews being particularly large in Slovenia. There is no statistically significant association between income group and retention probabilities in the majority of the countries. In those where there is an association, being in the richest fifth or second richest is associated with greater retention probabilities than individuals in the poorest fifth (with Estonia being an exception where the opposite is true). In virtually every country, men and women aged under 40 have lower retention probabilities (other things being equal) and men and women aged 65+ have higher retention probabilities.(Romania and Bulgaria are exceptions: elders have relatively low retention rates.) Individuals living in single or couple households with a head aged 60+ tend to have higher retention rates than childless single adults aged less than 60. Controlling for income group membership, the association between retention and belonging to a lone parent household disappears in most countries (with the exception of Iceland and the Netherlands).

Household head's labour market activity status is not significantly associated with retention in many countries. Where there is an association, it is typically the case that compared to individuals in work, individuals who are retired or inactive for other reasons are less likely to be lost to follow-up, or unemployed individuals are more likely to be lost. (It tends to be one association or the other, but not both.) Latvia is an exceptional case with relatively retention rates that are lower for the 'other inactive' group than for workers.

The complex association between the education level of the household head and retention rates found in the univariate analysis remains in the multivariate analysis. There is no association between education level and retention rates in about half the countries. Where there is an association, individuals with heads with education to post-secondary or tertiary level have larger retention rates than individuals with household heads with education to the primary level in some countries (mostly countries with relatively low overall national retention rates), but the reverse is true in other countries (mostly countries with relative high overall national retention rates).

In sum, there is substantial diversity in the rates at which individuals from EU-SILC Wave 1 samples are also found in the four-wave Balanced Samples from which persistent poverty rates can be calculated. There is differential attrition in terms of observable characteristics. The finding of diversity in retention rates was also reported by Behr, Bellgardt, and Rendtl (2005) and Watson (2003) in the ECHP, though specific results are difficult to compare because findings are summarised in different ways in the different studies (and there is no good 'one number' summary of the *amount* of differential attrition).

Behr, Bellgardt, and Rendtl report that 'there is a tendency for young people to exhibit a slightly higher probability of not responding' (2005: 503), but they also find that differences related to personal characteristics are small relative to those related to survey information (whether the household had moved; whether the interviewer had changed). In comparison, for EU-SILC, we find more clear cut differential attrition related to personal characteristics (though young people are also more likely to be lost to follow-up). In our EU-SILC longitudinal files, we do not have as detailed survey information as was available in the ECHP, but aspects of such information do matter, as the results relating to proxy interview status show.

3.5 How much retention is attributable to observable differences between individuals?

The multivariate regression framework that underpins Figure 4 can be used to provide a summary measure for each country of the extent to which variations in individuals' retention probabilities are attributable to differences in their observed characteristics. The measure we use is McFadden's Pseudo-R^2 which is defined as 1 – (*LL*/*LL*0), where *LL* is the log-likelihood for the fitted probit model of retention with covariates and *LL0* is the log-likelihood of a probit model containing an intercept term only. The higher the Pseudo-R^2 is, the greater the variation in retention probabilities that is explained by the predictors included in the regression. This measure makes intuitive sense, though it does have a weakness if used to make comparisons across countries. The issue is that differences in retention probabilities vary with unobserved characteristics as well as observed ones, but the probit model necessarily assumes that the variance of the residual error (which encapsulates all unobserved differences) is equal to one.[(10)] That is, in cross-national comparisons, comparisons based on the Pseudo-R^2 assume that the dispersion in unobservables is the same – which of course may not be the case (and it is not possible to check this).

The estimates of the Pseudo-R^2 are shown in Figure 5. (Countries are ordered as before, according to their overall national retention rates.) For most countries, the fraction of variation in retention probabilities that is explained by the observable covariates is between around 1 per cent and 8 per cent, with some indication that the proportion explained increases with the overall retention probability. The percentage explained is thus relatively low, but this is not uncommon. For example, Watson (2003) also reported a similar finding for the ECHP.

Figure 5: How much of the variation in retention probabilities is accounted for by observed predictors?
(McFadden pseudo-R^2)

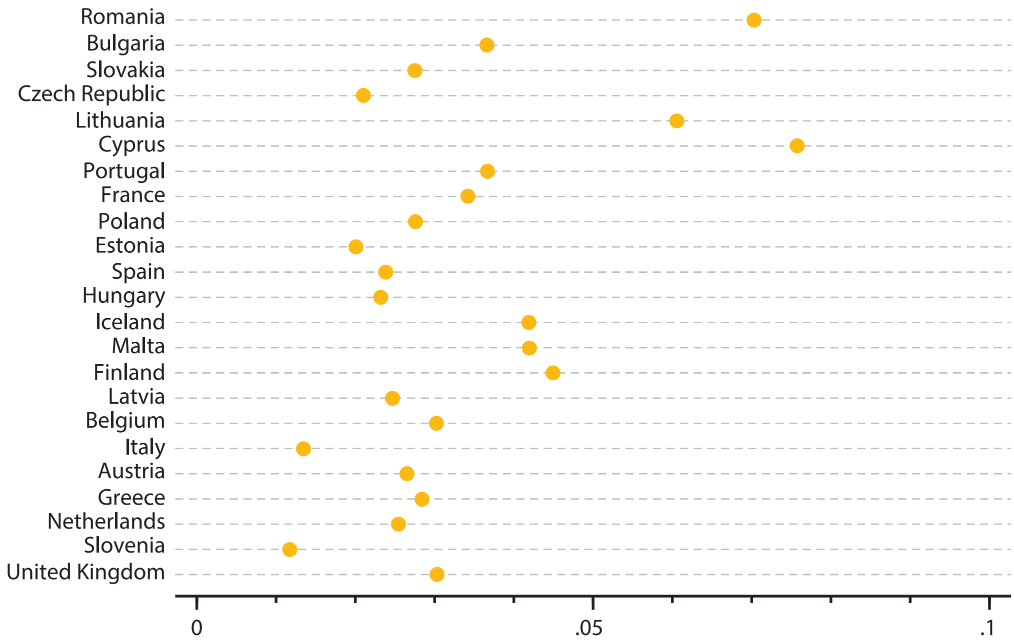

Note: The McFadden Pseudo-R^2 is 1 - (LL/LL0), where LL is the log-likelihood of the probit model of retention and LL0 is the log-likelihood of a constant-only model. The higher the Pseudo-R2 is, the greater the variation in retention probabilities that is explained by the predictors included in the regression (subject to the caveats explained in the text).

Source: Calculations from 2011 EU-SILC Longitudinal data; 2008 rotation group only.

(10) The unit variance normalisation is necessary to identify the parameters of the probit model. The problems arising when comparing regression coefficients across groups (e.g. countries) is well discussed by Long (2007). See also Williams (2009). If the outcome were a metric variable rather than a binary one, then the residual variance is estimable, and a partition of the total variance into components associated with observed and unobserved variables is possible. Watson (2003) uses the McFadden statistic for the same purpose as us but, because she pools the data all countries in her regression model, she implicitly assumes the same (common) variance.

3.6 Generating bespoke sample weights from retention regressions

Having the retention regression estimates means that we can construct our own bespoke longitudinal weights for counterfactual analysis of the four-wave Balanced Samples. The purpose of longitudinal weights in general is to adjust the four-wave Balanced Sample so that the reweighted Wave 1 covariate distributions of the sample is the same as the Wave 1 distribution of covariates in the full Wave 1 Sample. We construct our bespoke weights by multiplying the Wave 1 base weights, ω_i, of each observation *i* by the inverse of the retention probability predicted by the combination of the fitted regression parameters and the values of the predictor variables. So, if $r_i = \Phi(X_i b)$ is the predicted retention probability of observation *i* belonging to the four-wave Balanced Sample for a given country given standard Normal distribution $\Phi(.)$, vectors of characteristics X_i and fitted regression parameters *b*, individual *i*'s bespoke longitudinal weight is $w_i = \omega_i \times (1/r_i)$.

EU-SILC longitudinal weights are constructed by national statistical institutes in a similar though not identical fashion. They use cumulative year-on-year retention probabilities (rather than a four-year probability we have). They may include Wave 1 characteristics or perhaps more detailed fieldwork information that is not available in the public release files; and more flexible specifications for the regression equations used to predict retention probabilities; and there may be adjustment and calibration to the marginal distributions observed in full cross-section samples. In this paper we use only a basic specification, but it is one that can be implemented straightforwardly using the data in the scientific-use EU-SILC files that are available to us.

The greater the dispersion of subgroup attrition rates around a national average (as documented in sections 3.3 and 3.4 above), the greater the variance in a country's sample weights. This can also have an adverse impact on the precision of estimates of persistent poverty rates – a factor that we did not take into account earlier (in section 3.2) when showing the connection between standard errors and sample size. We assumed in our illustration that the design effect was constant across countries. However, the variance of the sample weights is one of the factors that influence the design effect, mediating the relationship between sample size, standard errors and the poverty rate. So, although application of sample weights may adjust for bias associated with differential attrition, it may come at a cost in terms of sampling variability when attrition is heavily differential and therefore sample weights have substantial variability.

In analysis that follows, we use our bespoke weights for some counterfactual exercises that can not be undertaken with the EU-SILC weights. We also use them in parallel with the EU-SILC longitudinal weights (so that we can check our specifications): see section 4.

3.7 How do patterns of differential attrition vary across countries?

In the first counterfactual exercise, we document how patterns of differential attrition differ across countries. The regression estimates shown earlier (Figure 4) alert us to the existence of differential attrition in every country, and different patterns across countries, but the implications of such differences are hard to assess in that metric. Moreover, differences in a country's retention probabilities depend on the distributions of individual characteristics (the relative prevalence of different types of individual), not only the regression coefficients. So, patterns of differential attrition can not be gauged directly from the latter.

These issues motivate an exercise in which we compare attrition patterns in terms of their effects on estimates of persistent poverty rates. That is, we consider questions such as: what would be the persistent poverty rate estimate for the UK if we applied the sampling weight adjustment required to correct for the attrition pattern of Romania (or Belgium or ...)? Into each country, we transplant every other country's attrition correction, thereby deriving 24 additional estimates that may be constrasted with the estimate corresponding to the country's own attrition pattern.

Imagine that retention rates are low among poor individuals in country *A*, equal to 1/3 (say), and assume that poverty status in the first wave is the only relevant determinant of attrition. By construction, sampling weight adjustments for country *A* would assign a large weight to the poor individuals who remain in the full balanced sample. That is, their weight would be multiplied by a factor 1/(1/3) = 3. Suppose that poor individuals in country *B* have a retention rate of 0.5, in which case the relevant adjustment for the non-attritors is 2. Our exercise consists in applying the adjustment factors of country *A* to the data of country *B* and therefore 'incorrectly' adjusting the sampling weights by a factor 3 instead of 2.

The extent to which the resulting estimates of persistent poverty rates is therefore indicative of how much sampling weight adjustments (and thereby differential attrition) differ across countries and how much it relates to calculations of persistent poverty. The advantage of this approach, which we have not seen employed before, is that it succinctly incorporates the impacts of both regression coefficients and distributions of characteristics. Moreover, this exercise is not subject to the problem cited earlier for coefficient comparisons across countries: predicted probabilities are not subject to this identification issue (Long 2007).

More specifically, the counterfactual exercise is undertaken by adjusting our bespoke longitudinal weights as follows. The weight for observation i is adjusted to country c's attrition pattern using the formula $w_{ci} = w_i \times r(X_i) / r_c(X_i)$, where $r(X_i)$ is the retention probability for an individual of characteristics X_i estimated from the probit regression model described above; $r_c(X_i)$ is the retention probability for an individual with the same characteristics in country c; w_i is the bespoke longitudinal weight and w_{ci} is the adjusted weight that is used to calculate the counterfactual persistent poverty rates.

The results of the exercises are shown in Figure 6. The various counterfactual persistent poverty estimates for each country are expressed relative to the estimate derived using the countries own attrition pattern. Transplant countries are identified by their two-letter acronyms in each country-specific row. If attrition patterns for transplant countries are similar to the reference country, estimates cluster around the vertical line at 1; if the attrition patterns differ substantially from the reference country, the counterfactual estimates are well away from this reference line. Put differently, a country with a distinctive pattern of attrition is either one with few counterfactual estimates close to its own estimate, or one that provides outlier counterfactual estimates for other countries.

Figure 6 shows that a significant minority of the countries have patterns of differential attrition that are distinctive. These are Belgium and the UK, and to a lesser extent, Austria, Finland, Greece, Hungary or Poland.

Consider, for example, the UK's attrition pattern. If its attrition adjustment were transplanted elsewhere, it would lead to substantial increases in persistent poverty rates in for example Bulgaria, Slovakia, Italy, the Netherlands, Slovenia (all with increases of between 20 per cent and 40 per cent) and especially Hungary (an increase of more than 60 per cent). Conversely, transplanting other countries' attrition adjustment onto the Wave 1 UK sample would lead to lower persistent poverty rates for the UK. (This is also true for Belgium.) This indicates that countries tend to retain more individuals with relatively large persistent poverty rates, and fewer individuals with relatively small persistent poverty rates.

Looking at the figure as a whole, there is a tendency for transplantation to lead to both greater counterfactual persistent poverty rates among some countries and to lower counterfactual rates among others, but there is no obvious pattern to this. For example, there is no systematic association with overall retention rates. It is also the case that countries that appeared distinctive in the earlier uni- and multivariate analysis are not necessarily the distinctive ones in the current analysis. This underlines the point that overall attrition overall can be seen as a combination of retention rates conditional on characteristics, and the relative prevalence of the groups of individuals with different combinations of characteristics.

Figure 6: Persistent poverty rates with transplanted attrition patterns
(ratio of observed country-rate)

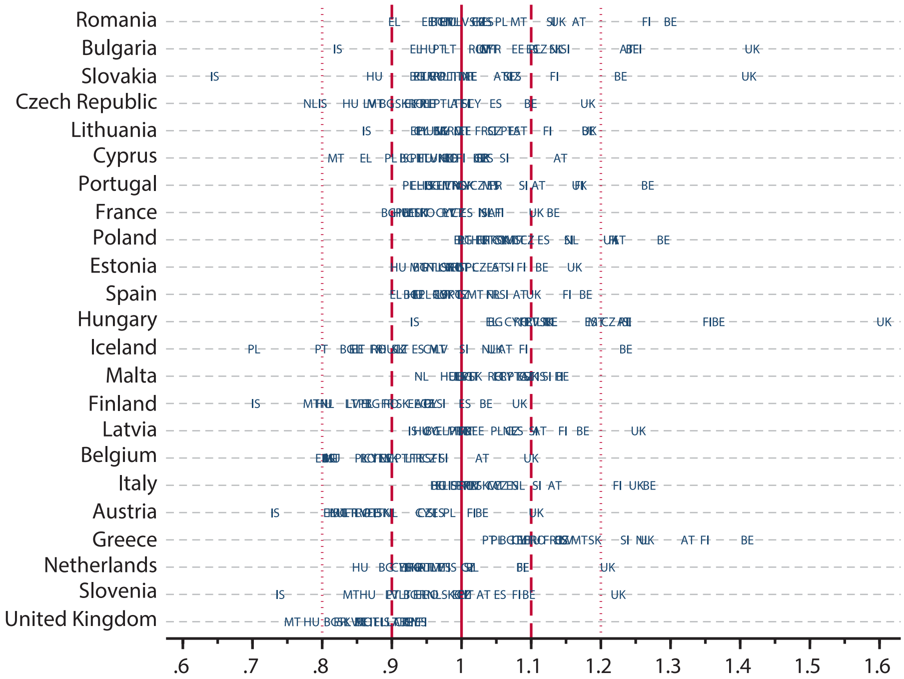

Note: Counterfactual persistent poverty rates for each country are obtained by applying attrition probabilities predicted from other countries onto wave 1 sample characteristics. All estimates use Wave 1 rb060 weights adjusted for (counterfactual) differential attrition using our bespoke retention regression model (see text). Each estimate is reported as the ratio of the corresponding estimate obtained from observed country-specific attrition.

Source: Calculations from 2011 EU-SILC Longitudinal data; 2008 rotation group only.

4. What effects does differential attrition have?

In this section, we begin by providing indirect evidence about potential bias in estimates resulting from differential attrition. Second, we use a parametric regression model of individuals' retention and persistent poverty probabilities to allow for the effects of unobservable factors on both outcomes. Third, we employ a novel non-parametric method to place bounds on persistent poverty rate estimates, where the bounds correspond to extreme assumptions about the poverty status of individuals who drop out of the initial W1 Sample.

4.1 Indirect evidence of attrition bias: comparisons of estimates of Wave 1 poverty rates

We follow Behr, Bellgardt, and Rendtl (2005) and assess the magnitude and potential impact of attrition by comparing our original base sample with the sample that remains after attrition. In our application, this means comparisons of statistics derived from the W1 Samples with the same statistics derived from the four-wave Balanced Samples. The benchmark statistic is the Wave 1 (2008) poverty rate. If there are differences in estimates, this suggests that the differences in the samples will also lead to bias in estimates of persistent poverty statistics (which can not be calculated for both samples, of course).(11) We refer to this as indirect evidence because it is not directly about the persistent poverty rate.

Figure 7, panel (a), contains four series of estimates of Wave 1 poverty rates for each country (countries are ranked as before). The 'All W1' sample consists of all individuals in the Wave 1 sample; there are two series calculated using our Balanced Samples but using different weights (the EU-SILC longitudinal weights and our bespoke ones). As an additional reference point, we show the estimates of the 2008 poverty calculated using the full cross-section sample (i.e. based on multiple cohorts rather than simply one) that are calculated and reported by Eurostat (2014).

In an ideal world, without random sampling variability, attrition or other data problems, all series would provide the same estimate for the poverty rate in 2008. Comparison of the Eurostat series (yellow filled circles) with the 'All W1' series (blue hollow diamonds) indicate that there are discrepancies associated with differences between full cross-sectional and longitudinal data files that we are unable to resolve using the data available to us. Although the differences in estimates are small for some countries (e.g. Slovakia, Bulgaria, France, Poland, Estonia, Hungary, Italy, Slovenia), there are differences that are relatively large for some other countries, nearly five percentage points in some cases (sometimes a negative difference, sometimes a positive one). See e.g. Romania, Iceland, Greece, the Netherlands, and the UK. There appears to be no systematic relationship with type of data collection or persistent poverty rate. Because these differences relate to aspects that we cannot control, we note the inconsistencies and pass on to the other comparisons. (Differences between full cross-sectional and longitudinal files are primarily related to sampling variability since the former are composed of about four times as many observations, but note also that the longitudinal files potentially have more non-response bias and older frames than 'All W1' estimates.)

Specifically, we are interested in the extent to which the longitudinally-weighted estimates from the four-wave Balanced Samples match the estimates from the 'All W1' samples (compare blue crosses and diamonds), and then the extent to which estimates using our bespoke weights match the estimates based on the Eurostat weights (compare yellow circles and blue crosses). These comparisons are easiest to make if one looks at Figure 7, panel (b), in which each longitudinal sample estimate is expressed as a ratio of the corresponding 'All W1' sample estimate. Longitudinal estimates that lie outside the boundaries demarcated by the vertical dashed lines differ by more than 10 per cent from the 'W1 sample' estimates. These cases are a signal that differential attrition is likely to lead to bias that is not fully accounted for by weighting. (The boundaries are sufficiently wide to allow for differences arising from sampling variation or differences in the underlying sample frame.)

The figure shows that, for 17 of the 23 countries, the longitudinal estimates based on the Eurostat weights are within 10 per cent of their full Wave 1 Sample counterparts. However, for 6 countries, the estimates are outside the boundaries, and hence there is indirect evidence that unaccounted-for differential attrition is possibly leading to bias. For three countries (the Czech Republic, Slovenia, and the UK), longitudinal sample estimates are between 80 per cent and 90 per cent of their corresponding Wave 1 sample estimate and, for three countries (Finland, Iceland), the longitudinal estimates are even smaller, less than 80 per cent of their corresponding Wave 1 sample estimate.

(11) The W1 Samples and the Balanced Samples are not representative of exactly the same target population. The latter excludes by definition people who leave the sample frame between wave 1 and wave 4 (through, death, out-migration, or move into non-private households). So differences in wave 1 estimates between the two samples can also reflect differences in target populations. However, over only four years, we expect this effect to be very small.

Figure 7: Estimates of Wave 1 (2008) poverty rate, by subsample and sample weight
(% and ratios)

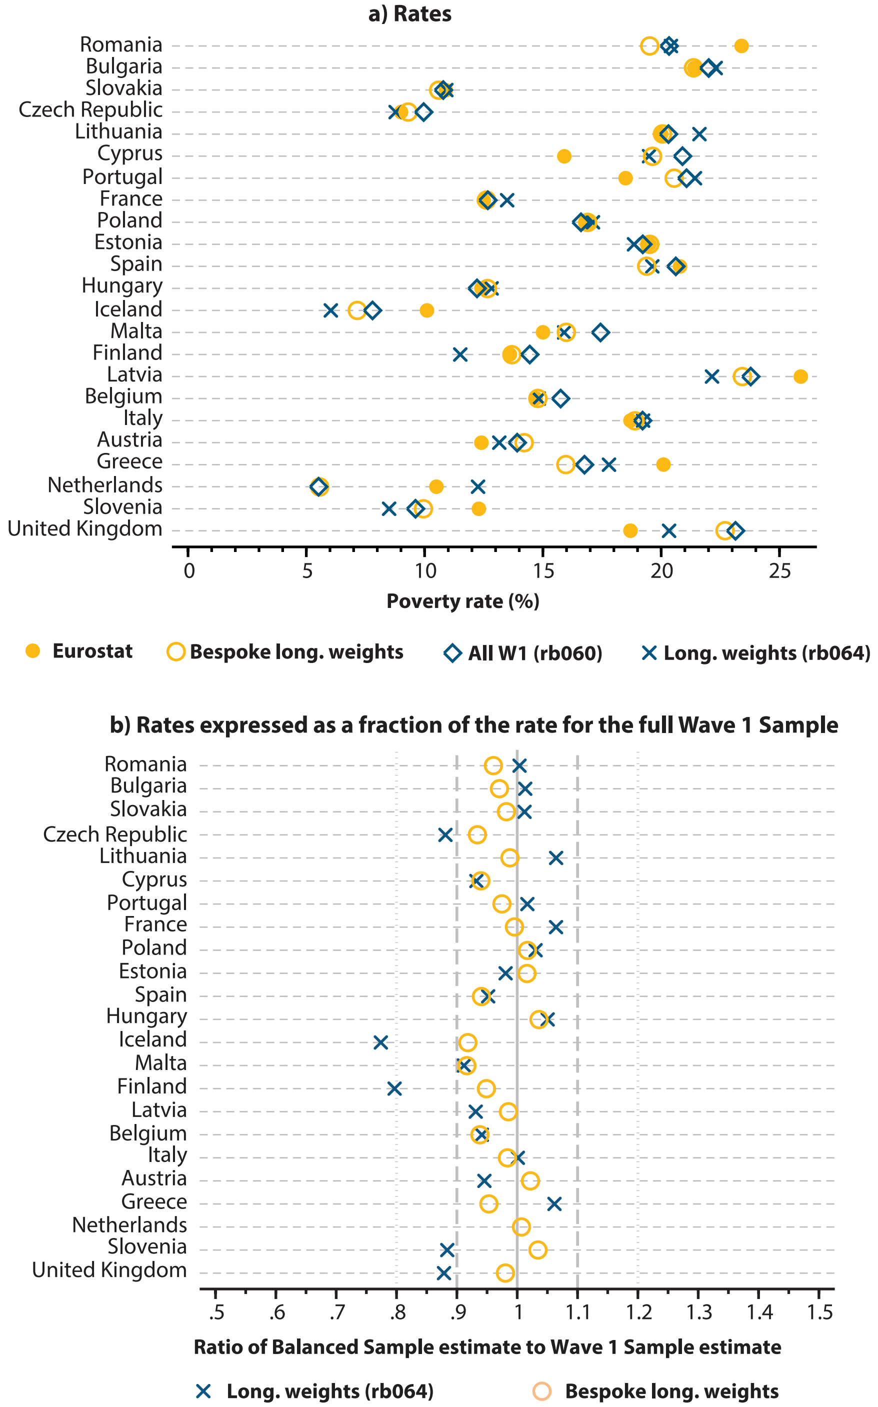

Note: All estimates are based on the balanced sample of full 4-year respondents, except 'All W1' which are based on all respondents at wave 1 (including subsequent attritors) and Eurostat's estimates based on cross-section data for 2008. The weighting of the samples are: 2008 base weights (rb060) for 'All W1', 2011 SILC longitudinal weights (rb064) and 2011 bespoke longitudinal weights for the Balanced Sample. Panel (b) shows estimates expressed as a fraction of the corresponding 'All W1' sample estimates. In the bottom panel, the estimate for NL is not shown as its estimate from the SILC longitudinal sample is an outlier (see text).

Source: Calculations from 2011 EU-SILC Longitudinal data; 2008 rotation group only. Eurostat estimates are from Eurostat (2014) and are computed using 2008 cross-section EU-SILC data.

Differences turn out to be large for the four register countries (Iceland, Finland, Slovenia and the Netherlands). However, again, the evidence about potential bias from differential attrition is less strong for these countries if our bespoke longitudinal weights are used rather than the Eurostat ones. With the former, the longitudinal estimates for all countries are within 10 per cent of their Wave 1 Sample counterpart.

The Netherlands stands out in this exercise. Its poverty rate estimate based on the balanced sample with EU-SILC longitudinal weights is more than twice the estimate obtained on the All Wave 1 sample and the estimate obtained from the balanced sample with our bespoke weights. (The relative difference is so large that it does not fit the horizontal scale of Figure 7 (b).) Eurostat's estimate is in between the two estimates although it is almost twice as large as our calculations. This suggests, first, that the other three cohorts forming the full cross-section dataset differ widely from the 2008–2011 cohort we focus on. Second, the EU-SILC weights appear to have been calibrated to fit a rich set of external distribution characteristics including some income distribution information. This casts serious doubt about the representativeness of analysis based on the Dutch longitudinal data, since, as we shall illustrate, the large adjustments to the longitudinal sample weights lead to wide sampling variability.

In sum, there is suggestive evidence of bias from differential attrition for a number of countries, but strong(er) conclusions are difficult to draw because of the inconsistencies across the different sets of estimates. We observe that our bespoke longitudinal weights generally do a good job of reproducing estimates derived using Eurostat longitudinal weights but, again, there a few marked differences.

4.2 Is attrition bias within the range of sampling variability?

Our ability to draw strong conclusions is also complicated by the fact that all estimates are subject to sampling variation, and this may overwhelm any differences in bias due to differential attrition. We illustrate this point in Figure 8, which shows estimates of persistent poverty rates from the four-wave Balanced Sample calculated using the Eurostat weights (crosses) and our bespoke weights (hollow diamonds), and their associated 95% confidence intervals.[12] For reference, also shown (using hollow circles) are the estimates published by Eurostat for the same period (see Eurostat 2014) and estimates obtained without any weights (hollow squares). Puzzlingly, there are some distinct differences between the published estimates and those based on the Eurostat longitudinal weights: the former are larger than the latter for Spain and Cyprus (around three percentage points in the latter case).

The main messages of Figure 8, however, are, first, that confidence intervals for persistence poverty rates calculated using Eurostat and our bespoke longitudinal weights overlap substantially in the vast majority of countries. There are some clear differences, to be sure, most notably for the Netherlands, but also for several other countries (such as Slovakia) – the countries for which we identified differences between the series in Figure 7.

The second lesson of Figure 8 is that confidence intervals for persistent poverty rates can be relatively wide. The ranges shown in the figure are of course similar to those suggested by Figure 2, but the lesson here is that the effects of differential attrition would have to be relatively large for differences to be significant in the statistical sense and, for example, to change the ranking of countries by persistent poverty rates.

The wide confidence interval for the Netherlands with the EU-SILC longitudinal weights connects to the discussion of Figure 7 and illustrates how large adjustments to sampling weights influences sampling variability. In light of these results, it is unclear whether the benefits in terms of bias reduction from calibrating sample weights to more reliable external information (here the cross-section data) are not offset by the increased sampling variability. A more detailed analysis in terms of mean squared error (which summarises both bias and variance in a single statistic) would be relevant here.

The cross-country ranking of persistent poverty rates shown in Figure 8 is broadly the same as the ranking that we reported in earlier work (Jenkins and Van Kerm 2014), though we should point out that the estimates are not directly comparable because the sets of countries differ (the 21 used in the earlier paper are not a subset of the 23 used here), and the EU-SILC data have been revised since the earlier study.

(12) Confidence intervals are calculated analytically using a linearization formula for the standard error of the mean (see e.g., Binder 1983): the persistent poverty rate is the sample average of the indicator function equal to 1 if an individual is poor at Wave 4 and at two or three of the previous years, and zero otherwise. Our confidence intervals are approximate because: (i) we ignore variability introduced by the estimation of the poverty line at each period (we take the cross-section poverty lines as fixed); (ii) we ignore the higher-level clustering of the sample at the primary sampling unit level as information on PSUs is unavailable in the EU-SILC longitudinal data for most countries; and (iii) we ignore the stratification of the sample since, again, the information is not available in the sample. The only survey design feature that we take into account is the clustering of individuals at the level of the household to which they belong at Wave 1.

Figure 8: Estimates of persistent poverty rates with different sampling weights
(%)

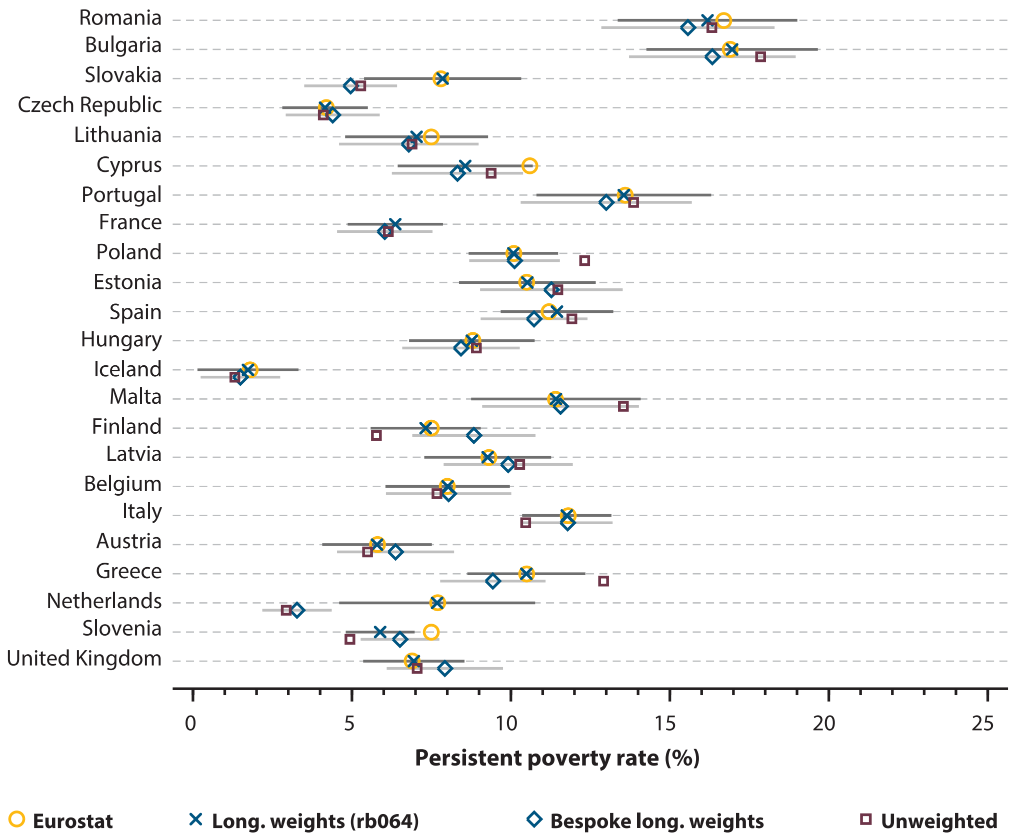

Source: Calculations from 2011 EU-SILC Longitudinal data (version 1, released August 2013). Eurostat estimates are from Eurostat (2014).

4.3 The effects of smaller sample size and differential attrition: the case of Romania

To further examine the twin impacts of attrition on estimate precision and bias, we undertake another simulation exercise, this time focusing on the case of Romania, the country with the largest national retention rate (over 90 per cent). The simulation involves artificially reducing the Romanian four-wave Balanced Sample by randomly deleting observations using deletion probabilities estimated from other countries. In other words, onto the Romanian sample, we project the attrition patterns of each other country. We can then recalculate the Romanian persistent poverty rate and its confidence interval in the simulated samples as if attrition were as in other countries (in terms of total amount and in its pattern). Differences between Romania's observed persistent poverty rate and associated confidence interval and the counterfactual counterparts provide information about the impacts of differential attrition and sample size reductions.

This simulation exercise is similar to the counterfactual reweighting shown in Figure 6. The key difference is that we now apply other countries' attrition on a reference 'best case scenario' in a more immediate way. For example, more than 90 percent of the initially poor remain in the sample in Romania, but the percentage is only about 45 percent in Slovenia (see Figure 3(a)). Our simulation applying Slovenian attrition to Romania data consists of randomly discarding half of the Romanians who were initially poor from their sample so that the simulated retention probability is the same as Slovenia's, and we then re-estimate measures of persistent poverty using the adjusted sample. In practice, deletion probabilities are calibrated to account for the full set of characteristics considered in Figure 3 and 4 (and not only poverty status in the illustrative example above). The exercise is conducted on the Romanian sample because the simulation can only involve *deleting* observations, and so we start with the sample with the highest retention rates. For the same reason, we cannot apply the simulation to Bulgaria, Latvia, Cyprus and France because for some subsets of characteristics, these countries have larger retention rates than Romania.

To isolate the effect of changing sample sizes from that of differential attrition, we add a reweighting step to the simulation. After dropping observations from the Romanian sample so that the revised sample size corresponds to the (lower) retention rates of another country, we re-estimated our bespoke longitudinal weights using the same regression model, but applied to the simulated sample. By construction, this step should undo the effect of differential attrition and reweight the sample back to the original Romanian full sample, except that the sample sizes are now smaller and some randomness is introduced by the (stochastic) simulation mechanism.

The results of the simulation exercise are shown in Figure 9. Panel (a) shows the results implied the adjustment to the longitudinal weights just described; panel (b) shows the results after the re-weighting step. In both panels, the estimated persistent poverty rate for Romania is shown at the top of the charts: 15 per cent, with a 95% confidence interval of [13.5%, 18.5%] shown by the two vertical lines.

Figure 9, panel (a), shows that, if Romania had the attrition pattern of other countries rather than its own, then its persistent poverty rate estimate would be substantially smaller. This is represented by the movement of the point estimates in the lower half of the figure to the left, beyond the lower bound of the confidence interval. However, one can also see, from panel (b) that the various counterfactual point estimates are brought back to the observed Romanian level once the weights have been re-adjusted: the estimates line up vertically below the Romanian estimate.

However, there remains the effect of having reduced the Romanian sample size to correspond with that of other countries, and this generally leads to wider confidence intervals. Recall that confidence interval width depends on the persistent poverty rate as well as its standard error which in turn depends on the persistent poverty rate and the sample size, so we do not expect to see the largest increase in confidence intervals in the rows corresponding to countries with the smallest decline in sample sizes due to attrition (Slovakia or the Czech Republic). The largest increase in confidence interval width turns out to be for countries with high attrition rates (Austria, the Netherlands, the UK; Figure 1(b)).

Figure 9: Persistent poverty rates for Romania with 'swapped' attrition patterns: effects on point estimates and 95% confidence intervals
(%)

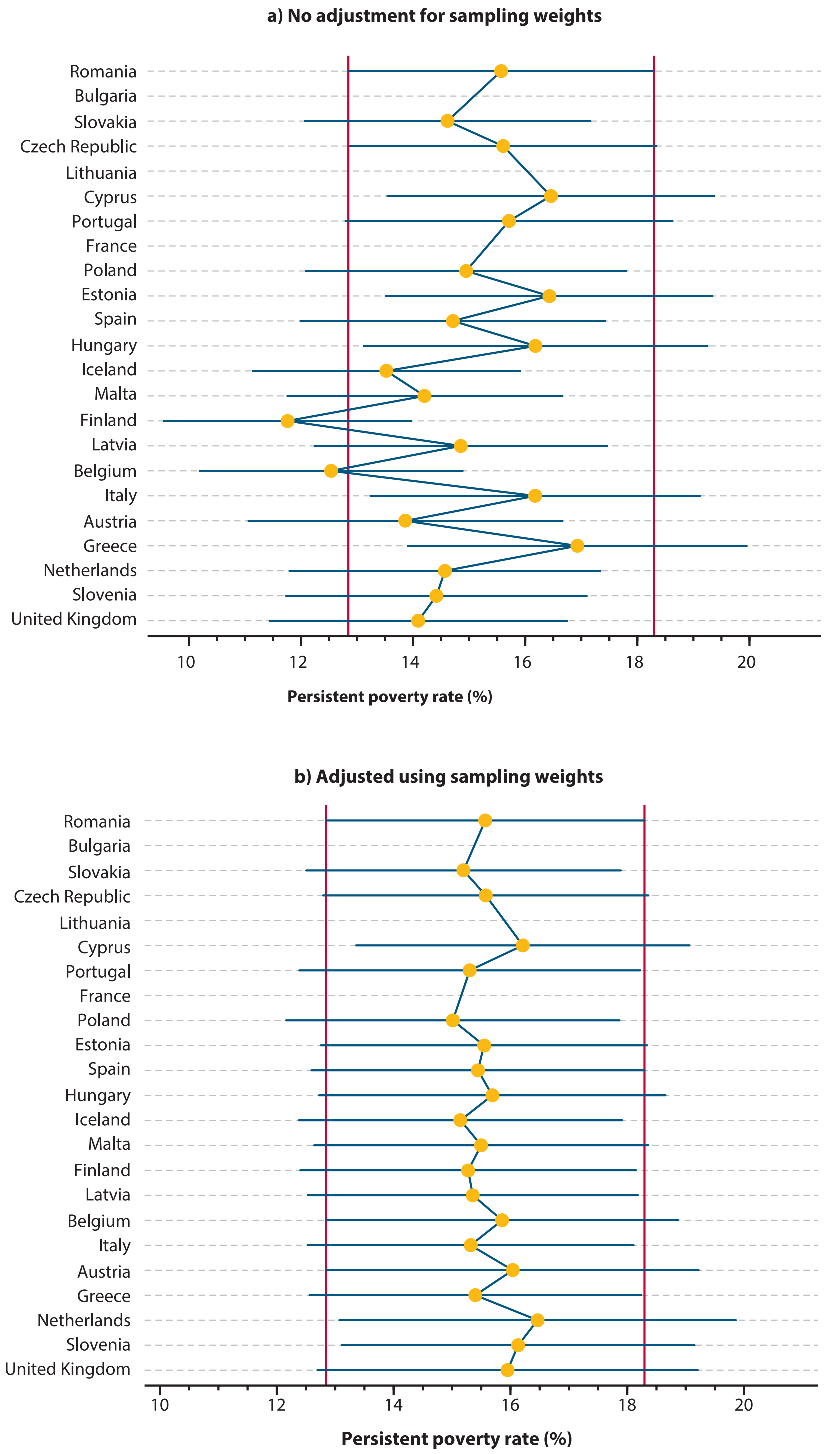

Note: Counterfactual persistent poverty rates for Romania are obtained by simulating higher attrition rates with the Romanian sample. For each other country c, each full respondent in Romania is dropped with a probability such that the overall retention rates and patterns correspond to that of country c. In the panel (a), persistent poverty rates are recomputed using each of the artificially-reduced samples without corrections to the sampling weights. In panel (b), our bespoke sampling weights are re-calibrated on the Wave 1 sample marginal distributions after the artificial elimination of observations.

Source: Calculations from 2011 EU-SILC Longitudinal data; 2008 rotation group only.

4.4 Allowing for unobservable effects using a parametric model

The strategy we have employed so far to take account of differential attrition is to suppose that its effects can be captured using observed variables, and sampling weights constructed from retention regressions and applied to counter the adverse effects. This is also the strategy employed by most data producers, including Eurostat.

There is a oft-stated concern that the coefficients in the retention regressions may be biased if there is a correlation between the unobservable factors that affect individuals' probabilities of sample retention and the probability of being persistently poor (even after conditioning on observed characteristics). One parametric strategy to address this issue is to fit a bivariate probit regression model with persistent poverty status and retention status as the two outcome variables and to allow for a potential correlation between the error terms of the two equations. Estimation of the model takes account of the fact that persistent poverty status is only observed for individuals who remain in the four-wave Balanced Sample (and is missing for individuals lost to follow-up). We would expect to find a negative correlation, i.e. unobservable factors that are associated with lower retention probabilities are likely to be associated with higher chances of being persistently poor.

These models are notoriously difficult to fit reliably. The main problem is to find variables that can be plausibly argued to affect retention probabilities but not persistent probabilities. Otherwise, identification of model parameters is entirely by functional form (including the assumption that the latent equation errors are bivariate normally distributed), and fragile. In fitting the model, we suppose that whether individuals had a proxy interview at Wave 1 is the variable excluded from the equation for persistent poverty status in each country. In the persistent poverty regression, the four dummies for income quintile groups are replaced by a single dummy for poverty status (in the initial time period) because in most countries almost none of those in the top quintile groups in the initial period are identified as persistently poor in the fourth wave.

The estimated coefficients from the persistent poverty regressions for each country are shown in the Appendix, together with the corresponding estimates derived from regressions that do not take account of correlated unobservables.

Figure 10 shows for each country the estimated correlation between the unobservables in the equations for retention and persistent poverty. A correlation significantly different from zero indicates that retention is correlated with persistent poverty even after controlling for the observed covariates. Estimates that are not statistically significant from zero are shown in grey. There are only 3 estimated correlations that differ significantly from zero and these are all very close to –1, which is a well-known signal of model fitting problems. (More generally, the model convergence was difficult for most countries.)

The implications for national estimates of persistent poverty rates of taking account of differential attrition related to unobservables can be examined by averaging the individual probabilities predicted from the model. (This exercise uses the point estimates of the correlations, regardless of their statistical significance.) We use the so-called marginal probabilities for the peristent poverty equation to generate the predictions (country by country), and then average over individuals within each country. We compare predictions based on jointly modelling persistent poverty and retention with persistent poverty predictions based on a model that ignores the potential corelation with retention (such predictions are identical to those obtained by direct calculation of the persistent poverty measures): see Figure 11.

For the few countries with 'credible' results (i.e. those with estimated correlations well away from the 1 and 1 extremes), this approach often leads to a slight reduction in the estimated persistent poverty rate relative to benchmark ignoring correlation in unobservables, but the correlations were not statistically significant in these cases. For the cases in which they were statistically significant but close to –1, the exercise leads to increases in estimates that are not credible. Overall, the parametric approach has little plausibility and so we move to our next approach.

Figure 10: Estimates of the correlation between the probabilities of retention and of persistent poverty
(%)

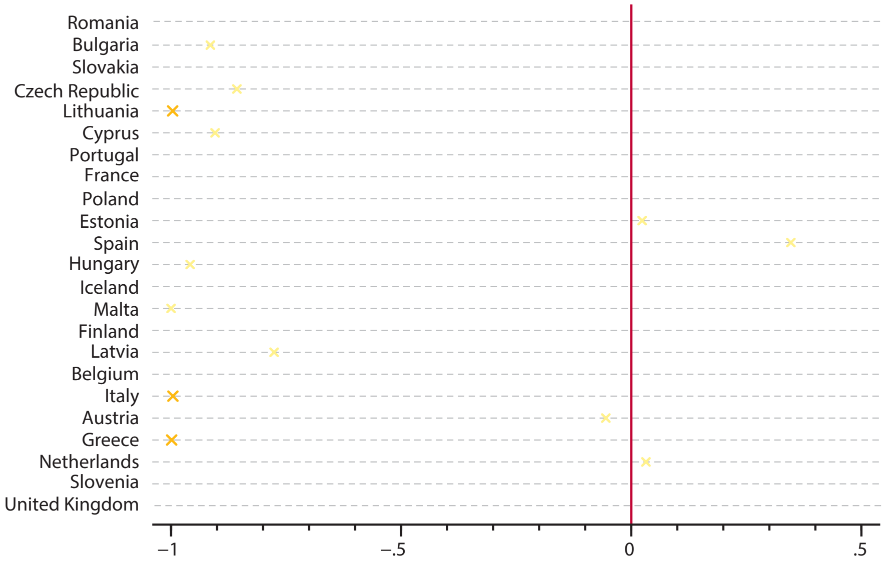

Note: The association between error terms in the two regression models is estimated in a transformed (arctan) metric to ensure the implied correlation lies between -1 and 1. Correlations for which the arctan-transformed parameters are not significantly different from zero are shown in the lighter shade. A negative correlation indicates that (conditional on the covariates), individuals with greater poverty persistence probabilities are less likely to be retained in the samples (i.e. have greater retention propensities).

Source: Calculations from 2011 EU-SILC Longitudinal data; 2008 rotation group only.

Figure 11: Persistent poverty rates predicted by joint models of poverty persistence and sample retention
(%)

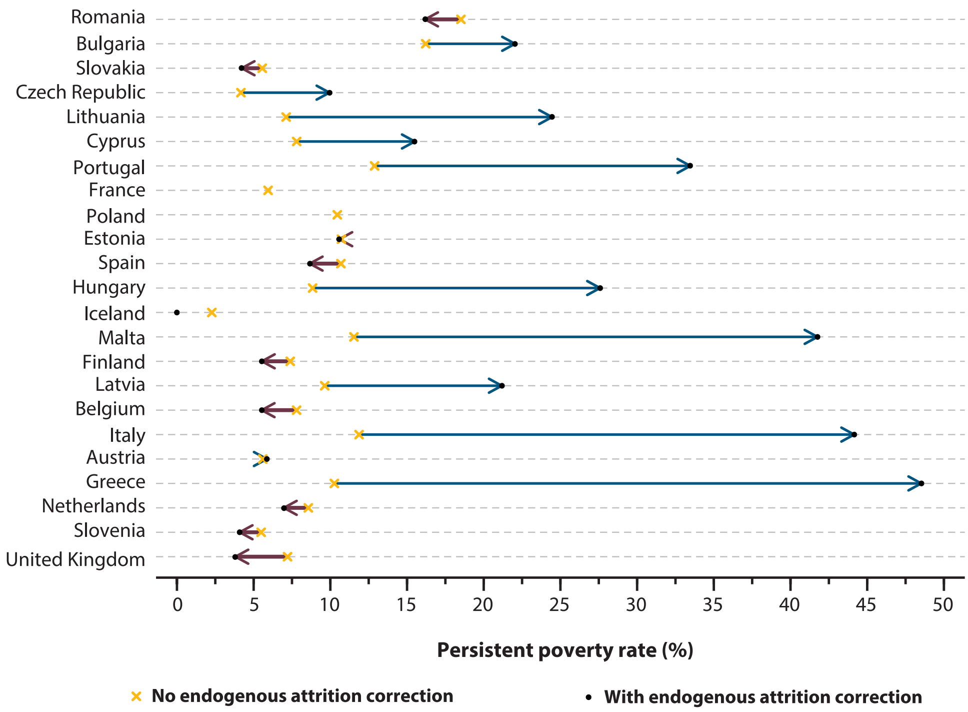

Note: The estimates are based on predicted probabilities derived from the bivariate model of poverty persistence and sample retention described in the text. The probabilities are averages over the balanced sample (with EU-SILC longitudinal weights rb064) of each individual's predicted probability of being persistently poor. These predicted probabilities are the marginal probabilities of being persistently poor (irrespectively of the retention status). In the model ignoring endogenous attrition, the estimates are very close to the persistent poverty rates observed in the raw data.

Source: Calculations from 2011 EU-SILC Longitudinal data; 2008 rotation group only.

4.5 A non-parametric bounding approach

The poverty status of attritors in their non-response years is unknown, by definition. Ignoring attrition altogether and estimating persistent poverty rates using the balanced subsample of respondents assumes that attrition is random and unrelated to one's poverty status. Applying longitudinal weights corrects for differential attrition to the extent that attrition is random *conditionally* on covariates observed prior to sample dropout. Although a weighting strategy is much better than ignoring attrition altogether, its validity hinges on the assumption that the poverty status of attritors does not systematically differ from the poverty status of respondents with identical characteristics in earlier interviews. In the previous subsection, we showed that parametric modelling approaches to this issue resulted in limited success. This leads us to a non-parametric bounding approach.

The motivating idea is that it is possible to calculate bounds within which a persistent poverty estimate for a country lie, whatever the underlying drivers of the attrition process. The approach of providing a range of estimates rather than a single point estimate is known as a 'partial identification' approach, and has been applied to date primarily to estimates of the current poverty rate, and using rather different types of information than we have available. See, for example, Nicoletti (2010), Nicoletti, Peracchi, and Foliano (2011), and references therein.

The way in which we calculate bounds with the data available in the longitudinal EU-SILC is as follows. It consists of different ways of imputing the poverty status of attritors in the years for which data about this status are not available.(13) The different imputation choices imply bounds on the true persistent persistent poverty estimate.

The first step consists in calculating, for each year (and each country), the number of respondents and non-respondents that should experience poverty in order to be consistent with the known poverty rate for that year. Let N denote the sample size at wave 1. Let p_t be the poverty rate in year t as calculated from the full cross-section sample available at time t, and which we derive from Eurostat (2014). To be consistent with this poverty rate, we should observe N_p=Np_t individuals in poverty at any year t in the full longitudinal sample as well. If we observe, among respondents at wave t, that $N_r p_t$ of them are poor, we can infer that $N_p - N_r p_t$ of the *non-respondents* are poor. Thus, imputation entails the allocation of these $N_p - N_r p_t$ poverty experiences to the set of N–N_{rt} non-respondents at wave t. (In the unlikely situation where there are fewer non-respondents than poverty experiences to allocate, we assume all non-respondents are poor, and therefore slightly underestimate poverty at time t.) To account for the fact that poverty rate derived from the full Wave 1 Sample may not be identical to the poverty rate derived from the full cross-section data at wave 1, we scale p_t by $(Np_1/N)/p_1$, i.e. the degree to which poverty is over- or underestimated at wave 1 in the full sample. We assume that the degree of over- or underestimation of poverty in the (imputed) longitudinal sample compared to the full cross-section is the same in each of the four years.

To calculate the upper bound on the persistent poverty rate, we allocate poverty experiences to non-respondents at year t on the basis of the sum of poverty experiences (observed or imputed) in year $s < t$. In year $t = 1$, all poverty statuses are observed and there is no imputation required. We then sequentially impute poverty statuses to individuals in the samples at year 2, then at year 3, and finally year 4. This means that at year 4, we allocate poverty experiences to all individuals who already have three years in poverty. If there are more poverty experiences to allocate than individuals with three years in poverty, we allocate them to individuals with two years in poverty in the previous three, etc. This allocation rule maximizes the persistent poverty rate under the constraint that the current poverty rates in the sample in each year are consistent with published cross-section current poverty rates.

To calculate the lower bound on the persistent poverty rate, the allocation rule is reversed. In each period, the poverty experiences are first allocated to individuals with no previous poverty exepriences in years $s < t$, then to individuals with only one poverty experience, and so on, until the number of poor people is reached.

In practice, the algorithms are adjusted to allow for differential sampling weights in the base year, and the cross-section poverty rates are consistent with weighted estimates of current poverty rates in each year.

The strength of this approach is that the bounds represent limits between which values for persistent poverty rates lie, regardless of any potential association between (unobserved) poverty (and poverty persistence) and attrition. Of course, the narrower the bounds, the better. The interval between bounds is a function of the attrition rate, the cross-section poverty rate, and the (unobserved) association between attrition and poverty status. Published estimates may lie outside the range shown by the bounds because our algorithm imposes a constraint on consistency with cross-section estimates.

(13) Since it is generally not possible to identify precisely the cause of non-response, we impute a poverty status to all individuals present at wave 1 whenever they do not provide data in any of the subsequent years. This is not optimal since some of these respondents have moved out of scope and not been eligible for interview (e.g. because of death or international migration) and should not contribute to calculations of the persistent poverty rate. We are confident that our assumption does not lead to significant bias since such individuals are likely to represent a small share of the total number of attritors. Lack of follow-up or refusal are the most prevalent forms of attrition.

The results of this bounding exercise are shown in Figure 12. It turns out that bounds are remarkably narrow for some countries, mostly those with relatively large retention rates, but their number is relatively small. Notable examples are Bulgaria, Portugal, and Slovakia. These are countries that have low overall attrition rates (they are towards the top of the figure), or relatively large retention rates for poor households, or relatively low cross-sectional (current) poverty rates, or a combination of all three features.

However, for most countries, the bounds are largely uninformative because they are so wide. They span an interval of up to 10 percentage points in the UK, Greece, Latvia, and Spain, for instance. Such wide ranges illustrate the inherent uncertainty in estimates when there is non-response, and the potential error associated with assuming that attrition is completely accounted for by observable (wave 1) characteristics.[14]

Figure 12: Estimated bounds on persistent poverty rates
(%)

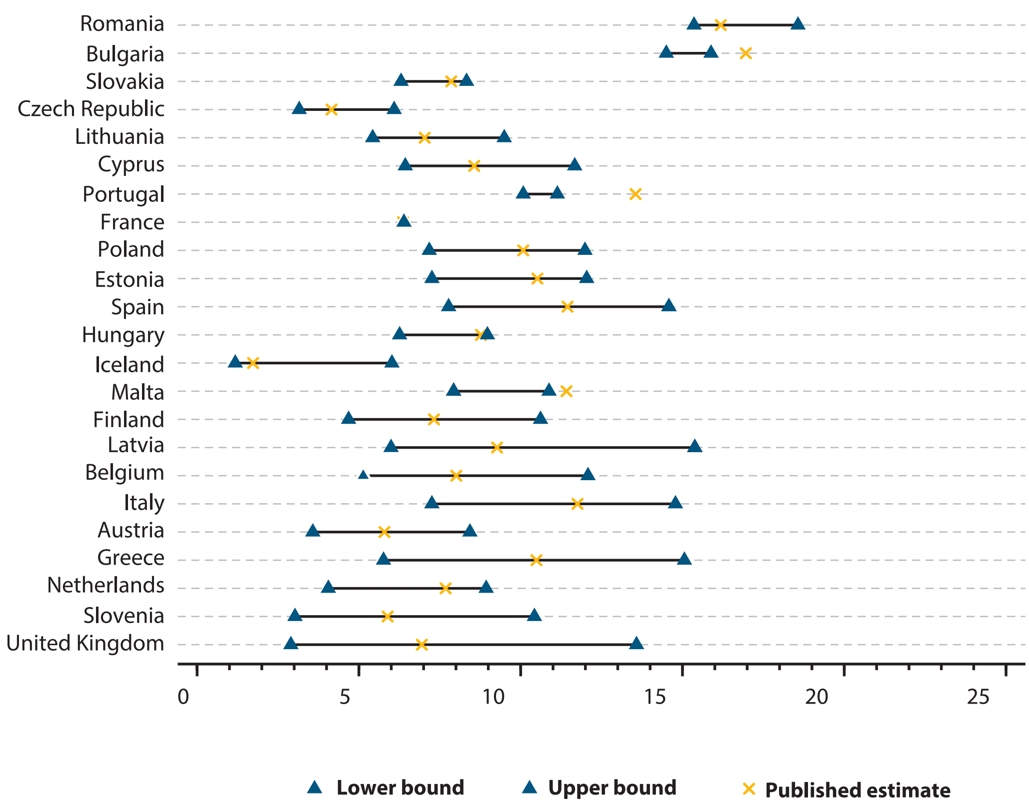

Note: The upper and lower bounds on persistent poverty rates are derived by imputing missing poverty status to attritors in different ways: see the main text for details.

Source: Calculations from 2011 EU-SILC Longitudinal data; 2008 rotation group only.

[14] There are no bounds presented for France. This is because there are not enough poverty experiences observed among respondents. Thus, even if everyone is imputed as being poor, we cannot not match the cross-sectional poverty rate that provides the relevant benchmark. We conjecture that may arise if there are large differentials in poverty rates between rotation groups and hence a differential also with the overall poverty rate.

5. Summary and conclusions

Rates of attrition from the four-year EU-SILC samples used to calculate persistent poverty rates vary substantially across Member States. The loss of sample size associated with attrition may lead to increases in standard errors for persistent poverty rates, and wider confidence intervals, that are sufficiently large – especially for population subgroups – that it is not possible to derive statistically robust conclusions about changes in persistent poverty rates over time or differences between subgroups.

There is substantial cross-national diversity in the characteristics of individuals lost to follow up. One key distinction stems from the different following rules applied by register and survey countries when households split. Overall, differential attrition abounds in the EU-SILC. We have shown that transplanting patterns of differential attrition from one country to another would lead to large changes in persistent poverty rates from observed rates.

We provide indirect evidence that application of longitudinal weights is essential yet it may not fully account for the effects of attrition, and that different assumptions about the poverty status of attritors lead to wide bounds in estimates of persistent poverty rates for most Member States. Thus, overall, we are less sanguine about the impact of attrition of EU-SILC-based estimates of persistent poverty than Watson (2003) was about the estimation of cross-sectional statistics using the ECHP.

We have been unable to pin down with confidence the effects of sample attrition on bias and precision in estimates of persistent poverty rates, but we have produced sufficient evidence to support a conclusion that researchers and data producers need to be mindful of these issues. Researchers analysing persistent poverty should at the least provide standard errors and confidence intervals for their estimates of rates, and their changes over time or differences between countries. (For EU-SILC applications, see e.g. Goedemé 2013.) Sampling variability is not identical to the uncertainty introduced by attrition, but accounting for the former should help remind readers of the effects of the latter.

National data collectors and Eurostat should continue their efforts to reduce loss to follow-up. This would be all the more important if a decision is made to extend the panel dimension of EU-SILC to more than four years. Extending the panel duration is attractive on substantive grounds, but minimizing attrition and minimizing cross-country differential in its patterns is essential to reap the benefits of this lengthening. Currently, the following rules are not implemented successfully in all the survey countries and, in the register countries, they are different by design. Our analysis has provided additional information about the groups at greatest risk of not providing income data for four years, and whom should therefore receive special attention. Reducing inconsistency across countries in the application of following rules (Iacovou and Lynn 2013) would also have payoffs for sample retention overall and reduce differential attrition. If more information about the details of the data collection process were available in the EU-SILC User Database, this might be used to derive better weights to account for attrition or to build more successful parametric models. We have also drawn attention to a number of apparent inconsistencies in estimates between cross-sectional and longitudinal components of the EU-SILC, and it would be good to have these resolved.

References

Behr, A., Bellgardt, E. and Rendtel, U. (2005), "Extent and determinants of the extent of panel attrition in the European Community Household Panel", ***European Sociological Review***, 21(5): 489-512.

Binder, D. A. (1983), On the variances of asymptotically normal estimators from complex surveys, *International Statistical Review*, 51(3): 279-292.

Eurostat (2010), "Description of target variables: cross-sectional and longitudinal. 2008 operation (Version January 2010)", Document EU-SILC 065, Eurostat, Luxembourg.

Eurostat (2014), Persistent at-risk-of-poverty rates, Online database, [Accessed: 20 April 2016]

Fitzgerald, J., Gottschalk, P., and Moffitt, R. (1998), "An analysis of sample attrition in panel data The Michigan Panel Study of Income Dynamics", *Journal of Human Resources*, 33(2): 300-344.

Goedemé, T. (2013), "How much confidence can we have in EU-SILC? Complex sample designs and the standard error of the Europe 2020 poverty indicators", *Social Indicators Research*, 110(1): 89-110.

Iacovou, M. and Lynn, P. (2013), "Implications of the EU-SILC following rules, and their implementation, for longitudinal analysis", *ISER Working Paper* n°2013-17, Institute for Social and Economic Research, University of Essex, Colchester.

Jäntti, M., Törmälehto, V.-M., and Marlier, E. (2013), *The use of registers in the context of EU–SILC: challenges and opportunities,* Eurostat Methodologies and working papers, Publications office of the European Union, Luxembourg.

Jenkins, S. P. and Van Kerm, P. (2014), "The relationship between EU indicators of persistent and current poverty", *Social Indicators Research*, 116(2): 611-638.

Lohmann, H. (2011), "Comparability of EU-SILC survey and register data: the relationship among employment, earnings and poverty", *Journal of European Social Policy*, 21(1): 37-54.

Long, J. S. (2007), *Comparing group effects in nonlinear models: statistical problems and substantive insights*, Presentation at Notre Dame University, April 18, 2007, Available at: http://www.indiana.edu/~jslsoc/files_research/rm4cldv/group_compare/long_group_nd_2007-04-16.pdf [Accessed: 20 April 2016]

Museux, J. M. (2006), *Weighting for EU-SILC*, Presentation at Methodological Workshop on 'Comparative EU-Statistics on Income and Living Conditions: Issues and Challenges', 6-8 November 2006, Helsinki.

Nicoletti, C. (2010), "Poverty analysis with missing data: alternative estimators compared", *Empirical Economics*, 38(1): 1-22.

Nicoletti, C., Peracchi, F., and Foliano, F. (2011), "Estimating income poverty in the presence of missing data and measurement error", *Journal of Business Economics and Statistics*, 29(1): 61-72.

Watson, D. (2003), "Sample attrition between waves 1 and 5 in the European Community Household Panel", *European Sociological Review*, 19(4): 361-378.

Williams, R. (2009), "Using heterogeneous choice models to compare logit and probit coefficients across groups", *Sociological Methods & Research*, 37(4): 531-559.

Appendix: Coefficient estimates from bivariate probit model with selection, by comparison with univariate probit model for persistent poverty

a) Constant term

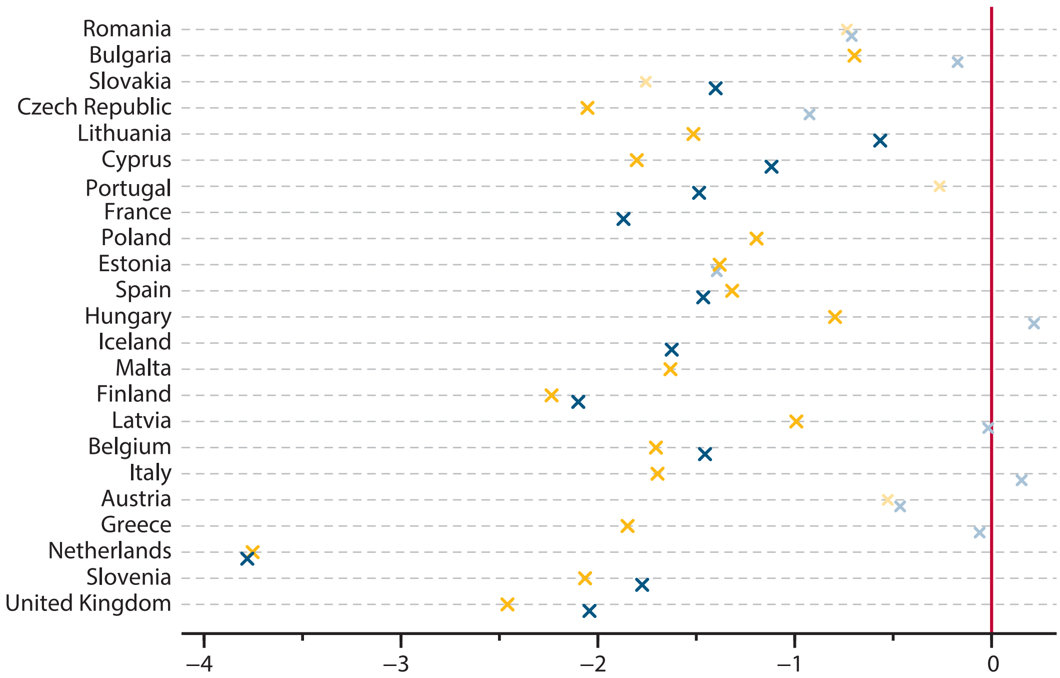

Constant term

b) Initial poverty status

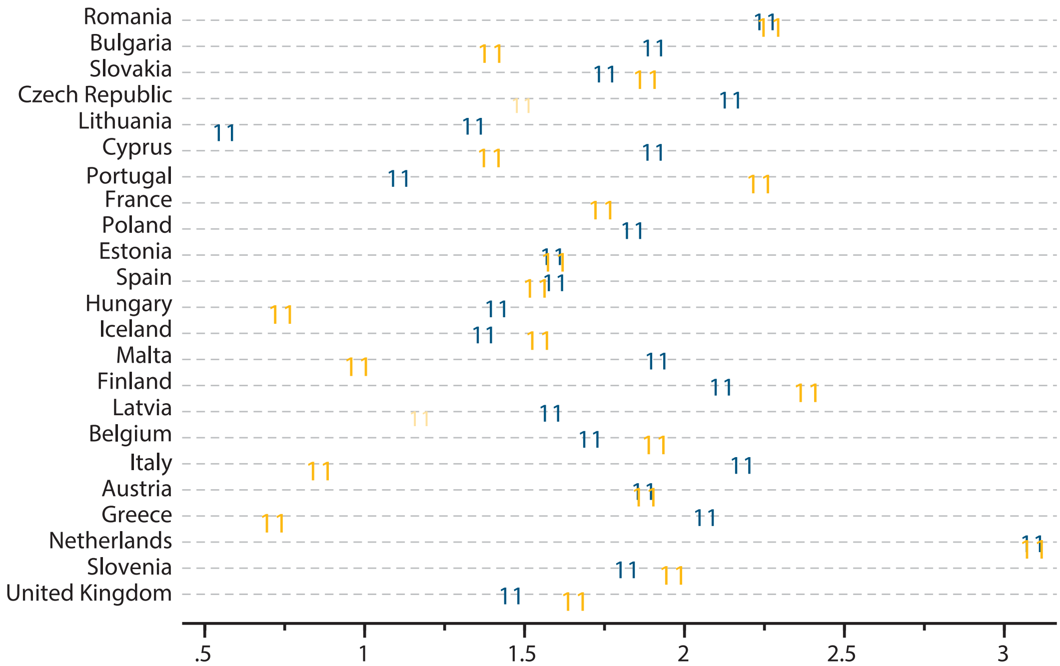

10 No (non–poor) **11 Yes (poor)**

c) Age-sex

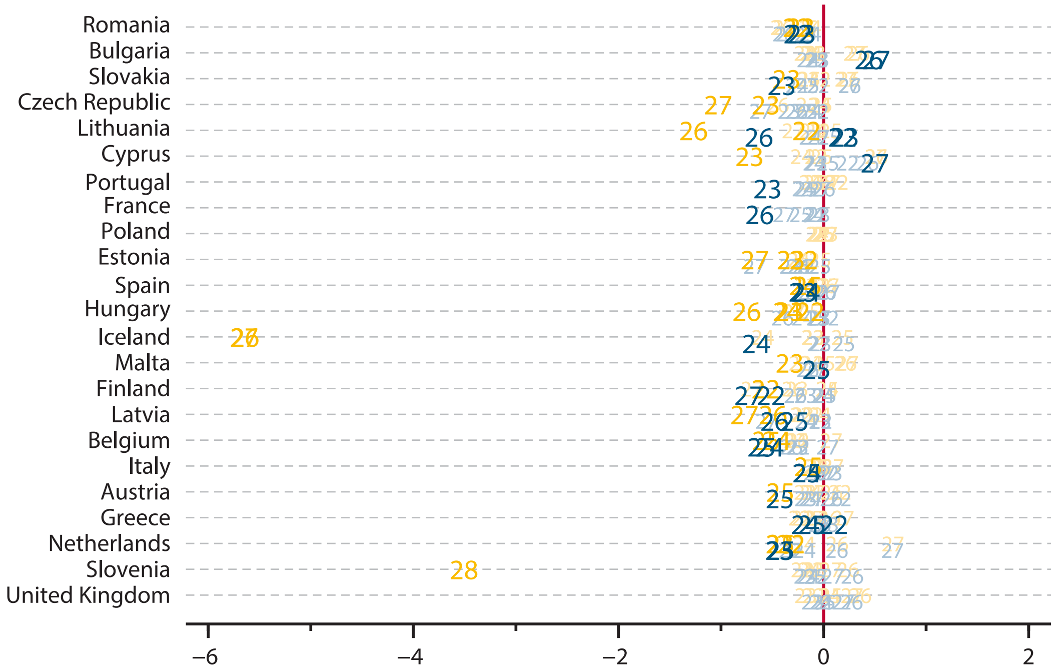

21 Under 1 **22 Woman aged <40** **23 Man aged <40** **24 Woman aged 40–64**

25 Man aged 40–64 **26 Woman aged 65+** **27 Man aged 65+** **28 Other/undefined**

d) Household type

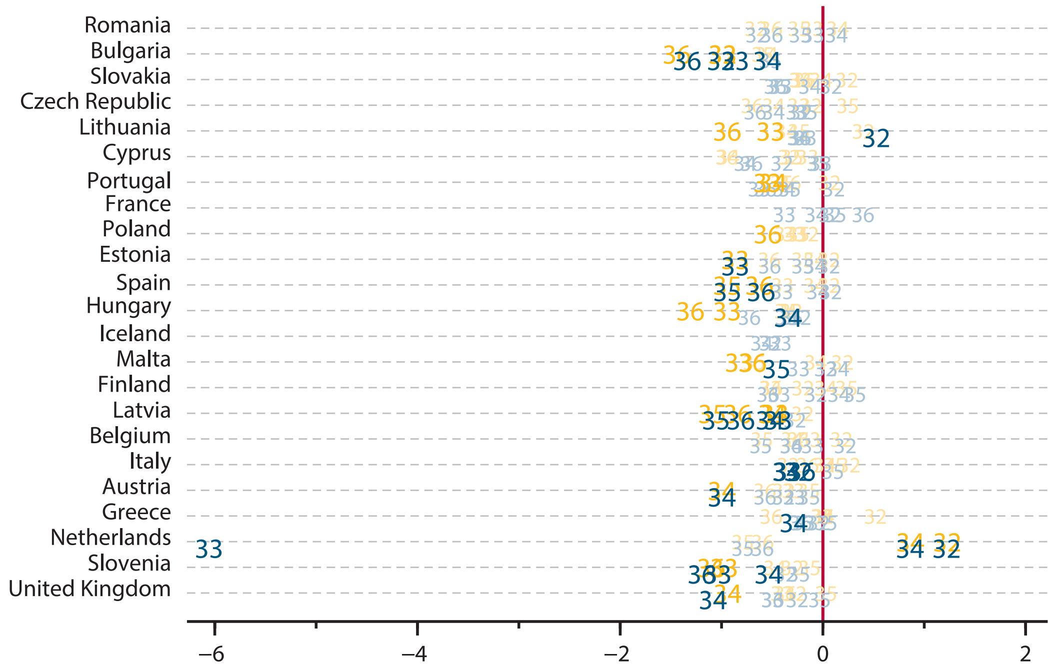

31 Single adult under 60, no kids **32 Single adult under 60 with kid** **33 Couple under 60, no kids**

34 Couple under 60, with kids **35 Single aged 60+** **36 Couple aged 60+**

e) Activity status of household head

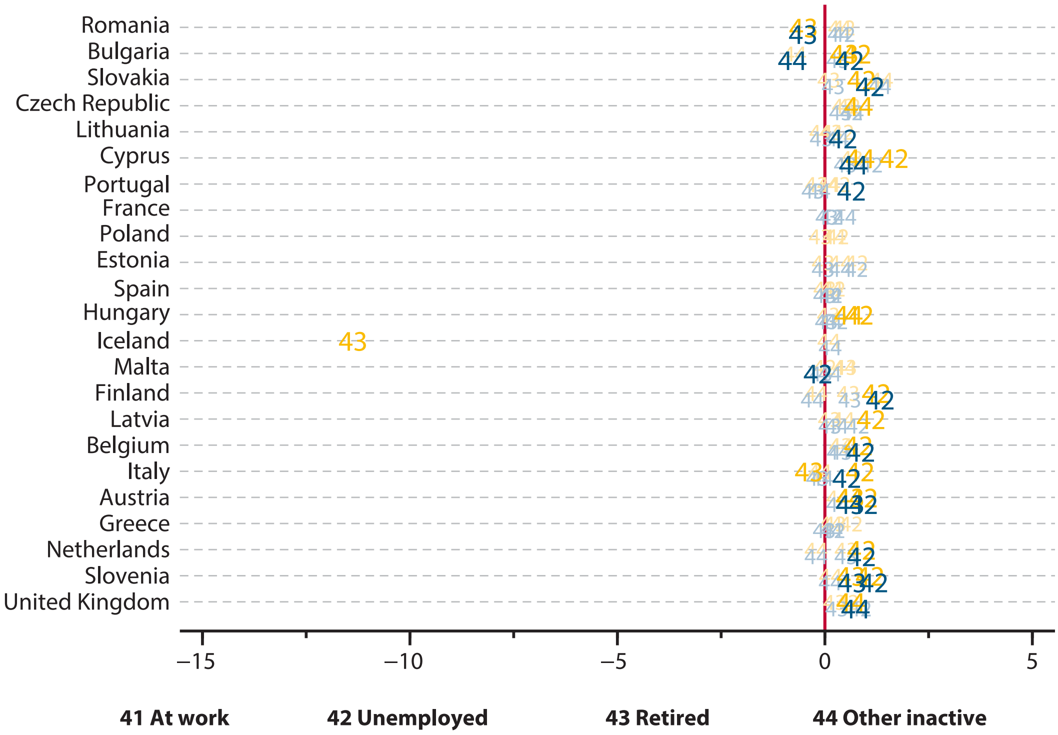

f) Education of household head

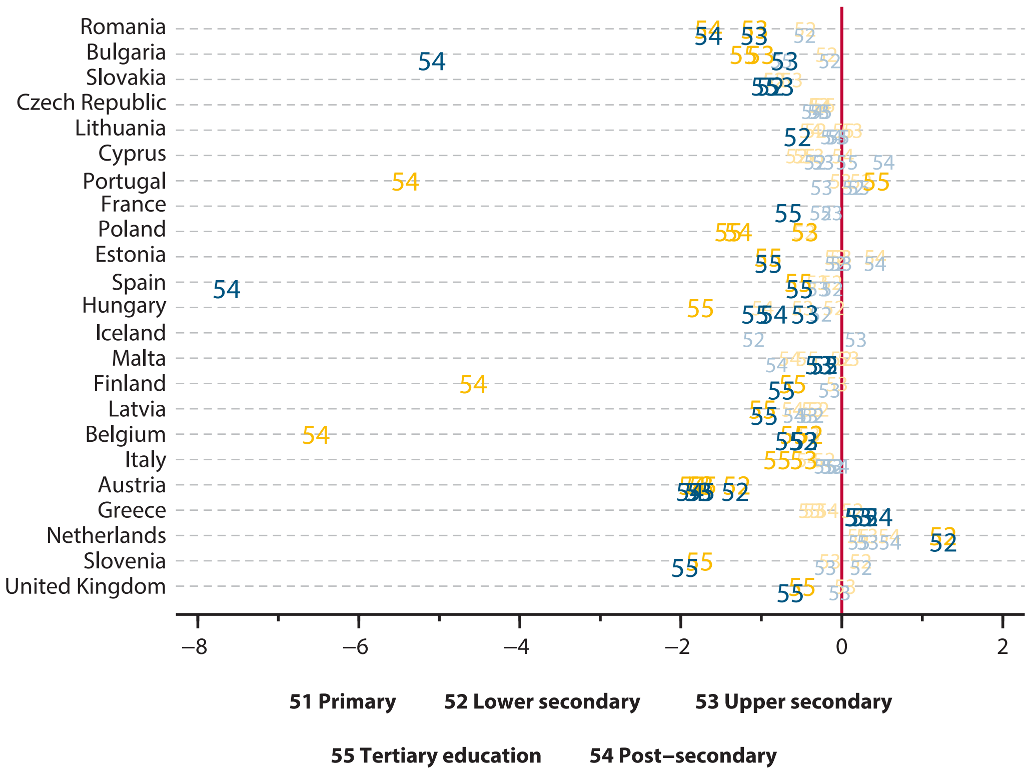

Source: Calculations from 2011 EU-SILC Longitudinal data; 2008 rotation group only.

HOW TO OBTAIN EU PUBLICATIONS

Free publications:

- one copy:
  via EU Bookshop (http://bookshop.europa.eu);
- more than one copy or posters/maps:
  from the European Union's representations (http://ec.europa.eu/represent_en.htm);
  from the delegations in non-EU countries (http://eeas.europa.eu/delegations/index_en.htm);
  by contacting the Europe Direct service (http://europa.eu/europedirect/index_en.htm) or
  calling 00 800 6 7 8 9 10 11 (freephone number from anywhere in the EU) (*).

(*) The information given is free, as are most calls (though some operators, phone boxes or hotels may charge you).

Priced publications:

- via EU Bookshop (http://bookshop.europa.eu).

KS-TC-16-025-EN-N

How does attrition affect estimates of persistent poverty rates? The case of European Union statistics on income and living conditions (EU-SILC)

Among the primary indicators of social inclusion is the persistent at risk of poverty rate, defined as the proportion of persons in a country who are at risk of income poverty in the current year and who were at risk of income poverty in at least two of the preceding three years. Evidence about poverty persistence is an important complement to information about poverty prevalence at a point in time. Estimates of persistent at risk of poverty rates are derived from the longitudinal component of EU SILC in which the fortunes of individuals are tracked over four consecutive years, in principle. In practice, not all of the individuals present in the first sample year provide four years of income data: there is attrition and estimates of persistent at risk of poverty measure may therefore not be reliable. Rates of attrition from the four-year EU-SILC samples used to calculate persistent poverty rates vary substantially across Member States, and there is also substantial cross-national diversity in the characteristics of individuals lost to follow-up. This working paper documents such patterns in detail and provides evidence that application of longitudinal weights does not fully account for the effects of attrition, and that different assumptions about the poverty status of attritors lead to wide bounds for estimates of persistent poverty rates for most Member States.

For more information
http://ec.europa.eu/eurostat/

Publications Office

ISBN 978-92-79-64143-5