SOEP The Socio-Economic Panel

1055 2019

SOEPpapers
on Multidisciplinary Panel Data Research

SOEP — The German Socio-Economic Panel at DIW Berlin 1055-2019

A practical guide for the computation of domain-level estimates with the Socio-Economic Panel (and other household surveys)

Natascha Hainbach, Christoph Halbmeier, Timo Schmid, Carsten Schröder

DIW SOEP

SOEPpapers on Multidisciplinary Panel Data Research at DIW Berlin

This series presents research findings based either directly on data from the German Socio-Economic Panel (SOEP) or using SOEP data as part of an internationally comparable data set (e.g. CNEF, ECHP, LIS, LWS, CHER/PACO). SOEP is a truly multidisciplinary household panel study covering a wide range of social and behavioral sciences: economics, sociology, psychology, survey methodology, econometrics and applied statistics, educational science, political science, public health, behavioral genetics, demography, geography, and sport science.

The decision to publish a submission in SOEPpapers is made by a board of editors chosen by the DIW Berlin to represent the wide range of disciplines covered by SOEP. There is no external referee process and papers are either accepted or rejected without revision. Papers appear in this series as works in progress and may also appear elsewhere. They often represent preliminary studies and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be requested from the author directly.

Any opinions expressed in this series are those of the author(s) and not those of DIW Berlin. Research disseminated by DIW Berlin may include views on public policy issues, but the institute itself takes no institutional policy positions.

The SOEPpapers are available at **http://www.diw.de/soeppapers**

Editors:
Jan **Goebel** (Spatial Economics)
Stefan **Liebig** (Sociology)
David **Richter** (Psychology)
Carsten **Schröder** (Public Economics)
Jürgen **Schupp** (Sociology)
Sabine **Zinn** (Statistics)

Conchita **D'Ambrosio** (Public Economics, DIW Research Fellow)
Denis **Gerstorf** (Psychology, DIW Research Fellow)
Katharina **Wrohlich** (Gender Studies)
Martin **Kroh** (Political Science, Survey Methodology)
Jörg-Peter **Schräpler** (Survey Methodology, DIW Research Fellow)
Thomas **Siedler** (Empirical Economics, DIW Research Fellow)
C. Katharina **Spieß** (Education and Family Economics)
Gert G. **Wagner** (Social Sciences)

ISSN: 1864-6689 (online)

German Socio-Economic Panel (SOEP)
DIW Berlin
Mohrenstrasse 58
10117 Berlin, Germany

Contact: soeppapers@diw.de

A practical guide for the computation of domain-level estimates with the Socio-Economic Panel (and other household surveys)

Natascha Hainbach[*], Christoph Halbmeier[†], Timo Schmid[‡], Carsten Schröder[§]

September 30, 2019

Abstract

There is a huge interest in deriving and comparing socio-economic indicators across societal groups and domains. The indicators are usually derived from population surveys like the German Socio-Economic Panel (SOEP) by direct estimation. Small sample sizes in the domains can limit the precision of these estimates. For example, while SOEP may be a suitable database for determining mean income in Germany, it is unclear whether this also applies to smaller domains (for example, women in Berlin). Here we show SOEP-based applications of Stata's `fayherriot` package (Halbmeier et al., 2019). This package implements the Fay-Herriot model (Fay and Herriot, 1979), a small-area estimation technique designed to improve the precision of domain-level direct estimates using domain-level covariates from auxiliary datasets.

Keywords: Disaggregated indicators, Small area estimation, Fay-Herriot model, Socio-Economic Panel

JEL Classifications: C13, C31, C51, C87

1 Introduction

An extensive empirical literature measures and compares socio-economic indicators (e.g. income, health, life satisfaction) across domains in a country. Such domains can be distinguished along factors like region of residence, gender, and birth cohort. The domain-level indicators are frequently derived from survey data by direct estimation. Along with this approach comes the central question whether domain sample sizes are sufficiently large to derive precise estimates. Institutions providing domain-level indicators usually require a minimum number of observations per domain or impose limits on the variability of the estimates (Eurostat, 2013; Tzavidis et al., 2018).

In Germany, many domain-level indicators rely on the Socio-Economic Panel (SOEP) (Goebel et al., 2019). One example is the "Glücksatlas," which presents life satisfaction

[*]SOEP at DIW Berlin, nhainbach@diw.de.
[†]SOEP at DIW Berlin, chalbmeier@diw.de.
[‡]Freie Universität Berlin, timo.schmid@fu-berlin.de.
[§]SOEP at DIW Berlin, cschroeder@diw.de.

indicators at the level of federal states (Krieg and Raffelhüschen, 2017).[1] Another example is OECD's project "Measuring regional and local well-being for policymaking," which provides estimators of inequalities within and between countries (Piacentini, 2014). Other examples assess disparities in health (Eibich and Ziebarth, 2014) and crime (Bug et al., 2015).

While the above list of SOEP-based domain-level works can be continued as desired, the issue of domain level sample sizes and precision of estimates, so far, has received little attention.[2] The idea of this short report is to alert SOEP users to the new statistical Stata package `fayherriot` for assessing and improving the precision of domain-level estimates (Halbmeier et al., 2019).[3] This package is easy-to-use and implements the area estimation technique suggested by Fay and Herriot (1979).

The Fay-Herriot (FH) model improves precision by combining the domain-level direct estimates from the survey data with auxiliary domain-level data from official registers or administrative data. The idea is that the auxiliary data is measured with greater precision and correlates with the outcome of interest. Thus, it contains additional information that is used to correct the direct estimates. Technically, the FH model uses the additional information to predict the outcome of interest with a linear regression model. Predictions and direct estimates are then combined into a weighted average, in such a way that more weight is given to the component that is estimated with larger precision. The result is a more precise domain-level indicator, the so-called empirical best linear unbiased prediction (EBLUP), which incorporates information from the survey and auxiliary data in an optimal way. In addition, the FH model estimates the mean squared error (MSE) of the EBLUP, which is typically used to assess its precision. Apart from valid and predictive auxiliary data, the FH model requires that the direct estimate from the survey should be a linear statistic such as an arithmetic mean, total, or share. Moreover, the model rests on the assumption of normal error terms.

The remainder of this note provides two SOEP-based applications of the `fayherriot` command to illustrate its functionalities step by step.

2 Empirical implementation

Initially one has to define the variable of interest from SOEP as well as the domain-level at which the analysis should take place. In our application, using SOEP data from version v34, we want to measure the average equivalent[4] yearly pre-tax income as well as the average satisfaction with one's home, both in 2015, at three regional levels, namely federal states, planning regions (Raumordnungsregionen) and districts (Landkreise). Note that the domain need not be a regional entity. For example, domains can also be defined along sociodemographic characteristics like age, household composition, or education level. The construction of the domains, however, requires mutual exclusiveness: no individual or household can be in more than one domain at the same time.

In a first step of the FH estimation, one has to calculate the direct estimates for

[1] See also Deckers et al. (2016) for a study on the relationship between satisfaction and regional prices.
[2] An exception is Eilers (2019).
[3] For R users see the package of Kreutzmann et al. (2019) and Molina and Marhuenda (2015).
[4] Here we use the OECD scale.

each domain, which in our case are the averages, and the sampling error variances. In case of the equivalent pre-tax income, averages are taken over all households in a domain because it is a variable measured at the household level, while averages of satisfaction with home are taken over all individuals. The sampling error variances reflect the imprecision of the direct estimates. We estimate the variances with the random group estimator to account for the survey sampling design as proposed in Rendtel (1995). Note that we use cross-sectional survey weights for these calculations.

Next, the FH model requires auxiliary statistics. These statistics should be at the same domain level, contain valid information, and be correlated with the SOEP-outcome of interest. These data cannot be from SOEP, but from some other reliable sources, collecting aggregated information for the relevant domain. Optimally these data come from registers or administrations. For our example, we retrieve additional data from the INKAR database, which regularly collects multiple indicators at different regional stages in Germany (Bundesinstitut für Bau-, Stadt-, und Raumforschung Bonn, 2018). For the mean equivalent pre-tax income, the auxiliary data contains the unemployment rate, the per-capita income tax revenue, and the share of people above 65 years. For the mean satisfaction with home, we use the mean household income, the share of apartments with just one or two rooms, and the population growth rate. Notice that our intention is not to select a comprehensive set of powerful explanatory variables, but rather a small set of variables that demonstrate the basic idea of the model and its application in Stata.

In many applications, some domains have no direct estimates, but auxiliary data is available. This is the case, for example, when some domains contain no or too few surveyed households to calculate/report direct estimates. We refer to these domains without direct estimates as out-of-sample domains, and to those with direct estimates as in-sample domains. The Fay-Herriot model also allows for calculating the EBLUP and MSE for out-of-sample domains. In such cases, the EBLUP simply corresponds to the linear prediction of the prediction model.

This said, the combination of the direct estimates with their sampling error variance and the auxiliary domain-level data from an additional reliable source are everything needed to make the `fayherriot` command work.

In the next subsections, we show the implementation of the `fayherriot` command with its pros and cons, also making a short comparison between the precision of the direct estimates and the FH estimates. For any further information regarding the `fayherriot` package and a comprehensive documentation of its features and options, see the package's help file.

Table 1: Number of regions and sample sizes for HH pre-tax income.

		Sample size distribution				
Regional division	Number of regions	Min	p10	p25	p50	Max
Federal states	16	117	149	452	634	3236
Planning regions	96	32	62	90	135	685
Districts	358	10	14	21	32	656

Note: Data are from SOEP v34. Own computations.

Table 2: Number of regions and sample sizes for individual satisfaction with home.

		Sample size distribution				
Regional division	Number of regions	Min	p10	p25	p50	Max
Federal states	16	187	251	730	1002	5577
Planning regions	96	49	98	154	231	1082
Districts	376	10	19	31	55	995

Note: Data are from SOEP v34. Own computations.

2.1 Sample size and descriptive statistics

For a quick overview, Tables 1 and 2 show the domain-specific numbers of SOEP households and individuals, respectively.[5] Apparently, at the district level, sample sizes are already small, challenging the precision of domain estimates of mean incomes and satisfaction.

The next two tables show the SOEP-based direct estimates – the mean equivalent pre-tax income and mean satisfaction with home – at the three aforementioned regional levels. Tables 3 and 4 provide the direct estimates. Incomes are in euros and satisfaction is measured on a scale ranging from 0 (lowest) to 10 (highest). In addition, the tables show a statistic for assessing the precision of the estimates, the coefficient of variation (CV), which is defined as the standard deviation of the mean divided by the mean (times 100). Note that there are many alternative criteria. The CV is used again later to compare the precisions.

Table 3: Summary statistics for equivalent HH pre-tax income.

Regional division	Min	p10	p25	p50	p75	p90	Max
(A) Mean equivalized household pre-tax income							
Federal states	13555	15187	16212	19369	24334	26258	28315
Planning regions	10891	14321	16793	20897	24017	27667	33343
Districts	4108	11507	15681	20933	26203	31466	110769
(B) Coefficient of variation							
Federal states	2.4	2.9	4.0	5.7	8.2	12.6	15.4
Planning regions	3.6	7.5	9.4	13.5	18.3	24.1	33.7
Districts	4.5	13.1	17.9	22.5	31.7	44.4	225.3

Note: Data are from SOEP v34. Own computations.

Mean incomes vary a lot between regions and the variation increases markedly from

[5]Note that we dropped regions with fewer than 10 observations for confidentiality.

Table 4: Summary statistics for individual satisfaction with home.

Regional division	Min	p10	p25	p50	p75	p90	Max
(A) Mean individual satisfaction with home							
Federal states	7.1	7.3	7.5	7.5	7.6	7.8	7.8
Planning regions	6.7	7.0	7.4	7.6	7.9	8.0	8.7
Districts	2.9	6.6	7.2	7.7	8.1	8.4	9.5
(B) Coefficient of variation							
Federal states	0.6	0.6	1.0	1.4	1.9	3.1	4.6
Planning regions	0.8	1.7	2.2	2.8	3.6	5.0	13.1
Districts	1.1	2.8	3.7	5.1	8.4	12.6	70.0

Note: Data are from SOEP v34. Own computations.

federal states over planning regions to districts. For instance, at federal states level, equivalent pre-tax income ranges from 13,555 to 28,315 euros, but at district level from 4,108 to 110,769 euros. For mean satisfaction, the pattern is similar: Mean satisfaction varies in a narrow interval from 7.1 to 7.8 at the federal level, but in a wide interval from 2.9 to 9.5 at the district level.

Thus, the point estimates show higher variability as the number of domains increases. The same holds for the coefficients of variation. As an example, according to Statistics Canada (2013), precision is too low when the CV is above 16.5 percent. Compliant with this threshold, estimates at federal states level are precise enough, but this is not true for estimates at the levels of planning regions and, in particular, districts. For example, CV for income goes up to 33.7 at planning regions level and exceeds 225 percent at district level. The variation for satisfaction is always smaller,[6] but also exceeds the value of 16.5 at district level. The next step is to explore whether `fayherriot` can help improve the precision of the domain estimates.

Table 5: Overview of data structure for HH income.

Domain	ID	income	direct variance	unemployment	income tax	share65	N
Planning regions	1	16003.39	3913200	7.5	372.1	21.7	119
	2	16734.15	8376093	7.1	332.1	22.8	167
	(...)	(...)	(...)	(...)	(...)	(...)	(...)
	96	18633.36	2.03e+07	5.5	249.9	24.8	109

Note: Highlighted columns in gray show the auxiliary variables from register.

2.2 Fay-Herriot estimation

The following example shows how to use the `fayherriot` command to estimate the FH model on the direct estimates of average pre-tax household income (`income`) as a function of the unemployment rate (`unemployment`), the per-capita income tax revenue (`incometax`), and the share of elderly people (`share65`) at the level of planning regions.

The command requires a combined dataset that contains the domain-level indicators

[6]Due to the short scale of the variable and the tendency of individuals to avoid extreme values, most values lie between 6 and 8.

from SOEP and auxiliary covariates from additional databases. Table 5 provides an exemplary overview of the data structure we use for the FH estimation of average pre-tax income at the level of planning regions. Columns 1 and 2 show the domain level with the respective ID. Column 3 displays the direct estimate of average pre-tax income (`income`) and column 4 the direct estimates of the sampling error variance (`directvariance`). Both are estimated with only SOEP data. Columns 5 to 7 contain the auxiliary covariates from register data, in our case from INKAR (Bundesinstitut für Bau-, Stadt-, und Raumforschung Bonn, 2018).

The command works similar to the normal regression command:

```
. fayherriot income unemployment incometax share65, ///
> variance(directvariance) gamma nolog

Sigma2_u estimation method:     reml            N in sample      =         96
Transformation of depvar:       none            N out of sample  =          0
EBLUP and MSE bias correction:  none            Sigma2_u         =  3.022e+06
                                                Adj R-squared    =     0.5804
                                                FH R-squared     =     0.8218

------------------- Gamma -------------------
  Min       5%      Median      95%      Max
0.0621    0.0933    0.2961    0.6106   0.8064

------------------------------------------------------------------------------
      income |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
unemployment |  -1.763695   171.2808    -0.01   0.992    -337.4679    333.9405
   incometax |   36.81214   4.934916     7.46   0.000     27.13988     46.4844
     share65 |   3.591693   225.5391     0.02   0.987    -438.4567    445.6401
       _cons |   5762.754   6438.612     0.90   0.371    -6856.695     18382.2
------------------------------------------------------------------------------

Shapiro-Wilk test for normality:
 Residuals e (standardized)  V =    1.217  p-value = 0.332
 Random effects u            V =    1.240  p-value = 0.317
------------------------------------------------------------------------------
```

The output table summarizes the prediction model. We see that the full set of 96 regions is used for estimation. The `FH R-squared` has a high value, 0.83, indicating that the prediction model has a good fit. The p-values of the beta coefficients show that, in particular, the `incometax` variable correlates significantly with the average pre-tax income. In addition, the Shapiro-Wilk test for normality does not reject the normality of residuals or random effects. Therefore, the model assumptions are not violated. Because we specify the `gamma` option, the table also shows summary statistics of the weighting factor gamma, which weights the predictions and direct estimates to calculate the EBLUP. Low values mean that little weight is put on the direct estimate and more weight on the model prediction. In the example, the median gamma value is 0.2958, indicating that for fifty percent of the observations, gamma is 0.2958 or smaller.

Now we show an example for the satisfaction with home variable (`homesatis`) at the district level. Here we decide to use the average household income (`HHincome`), the ratio of apartments with just one or two rooms (`lessrooms`), and the population growth rate (`popgrowth`) as auxiliary explanatory variables at regional level:

```
. use dataDistricts_homesatis.dta, clear

.
. fayherriot homesatis HHincome lessrooms popgrowth, ///
> variance(directvariance) gamma nolog

Sigma2_u estimation method:     reml            N in sample      =        376
Transformation of depvar:       none            N out of sample  =         26
EBLUP and MSE bias correction:  none            Sigma2_u         =     0.2375
                                                Adj R-squared    =     0.0673
                                                FH R-squared     =     0.1199

------------------- Gamma -------------------
  Min       5%      Median      95%       Max
0.0385    0.1389    0.6016    0.8621    0.9677

------------------------------------------------------------------------------
   homesatis |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
    HHincome |   .2660466   .1945511     1.37   0.171    -.1152665    .6473598
   lessrooms |   -.032005   .0074676    -4.29   0.000    -.0466413   -.0173687
   popgrowth |  -.0286521   .0175235    -1.64   0.102    -.0629974    .0056933
       _cons |   7.640313    .362757    21.06   0.000     6.929322    8.351303
------------------------------------------------------------------------------

Shapiro-Wilk test for normality:
 Residuals e (standardized)  V =    20.836  p-value = 0.000
 Random effects u            V =     6.989  p-value = 0.000
------------------------------------------------------------------------------
```

In this example, 26 districts are out of sample and not included in the prediction model because they have a missing value in `homesatis`. In contrast to the estimation for planning regions, the Shapiro-Wilk test rejects the normality of residuals and random effects, leading to the violation of the model assumptions. In addition, the `FH R-squared` measure gets very small (0.12). Hence, the FH model does not perform as well as for income at planning regions level.

What can be done if the normality assumptions are violated? First, one can transform the outcome variable into logs.[7] The `fayherriot` command has a `logarithm` option that ensures that EBLUPs and MSEs are transformed back to the original scale correctly and without bias. Second, one can exclude individuals or households that are outliers before calculating the direct estimates for each domain. Because as sample sizes get smaller, the more disaggregated the domains become, outliers can have a large impact on the direct estimates. Third, one can look for other auxiliary variables that explain the regional distribution of the outcome better. In some cases, however, one has to refer to other models that rely on a) other distributional assumptions or b) more flexible transformations (Rojas-Perilla et al., 2019; Fabrizi and Trivisano, 2016; Sugasawa and Kubokawa, 2015).

In case of `homesatis`, the model with logarithmic transformation still violates the normality assumptions. Because of a pre-defined range of the variable, exclusion of outliers is also not an option. For these reasons, we refrain from providing results for `homesatis` on the district level.

[7] Note that the sampling error variance also has to be transformed as shown in Halbmeier et al. (2019).

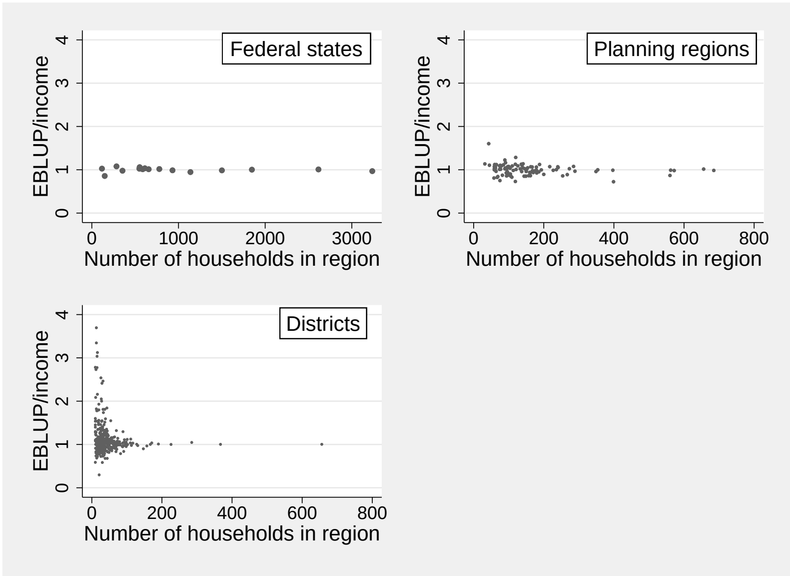

Note: Only in-sample domains are plotted. Data are from SOEP v34. Own computations.

Figure 1: Comparison of EBLUP and direct estimates for equivalent HH pre-tax income.

2.3 Comparison

Now we compare the direct estimates with the Fay-Herriot predictions (EBLUP). To calculate the EBLUP, we use the post-estimation `predict` command.[8] To assess the precision of the EBLUP, `predict` also allows for directly calculating the coefficient of variation based on the MSE. Figures 1 and 2 show the ratios of the EBLUPs to the direct estimates as a function of the domain-level sample sizes for the income and satisfaction variable, respectively. As expected, the lower the regional level, the more adjustments that must be made, due to smaller sample sizes. For the federal states, the direct estimates and the FH estimates are quite similar, whereas there are huge differences at district level.

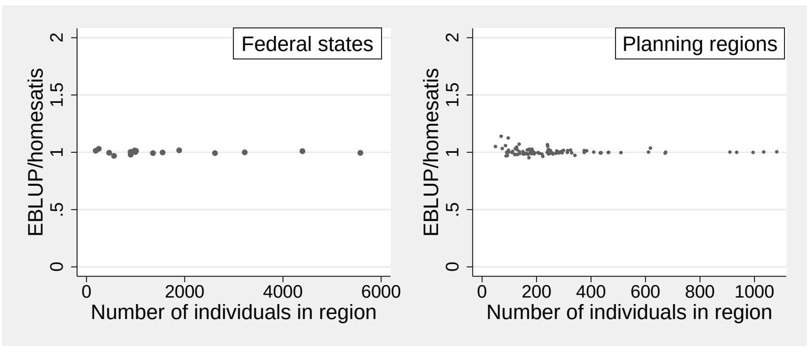

Note: Only in-sample domains are plotted. Data are from SOEP v34. Own coputations.

Figure 2: Comparison of EBLUP and direct estimates for individual satisfaction with home.

Figures 3 and 4 present boxplots of the coefficients of variation for the direct and FH estimates. It becomes apparent that many CVs of the direct estimates of income at the planning region or district levels are above the threshold of 16.5 percent. In contrast,

[8]There exists another way by specifying the `eblup(varname)` and `mse(varname)` option in the `fayherriot` command directly. For details, see the description of the `fayherriot` package.

the CVs of the FH estimates are significantly lower and mostly under the threshold level. Therefore, the results of the Fay-Herriot model are more precise than the direct estimates. For satisfaction, the FH model also reduces the CVs, but here the direct estimates already have CVs under the threshold of 16.5. Therefore, in both specifications, the FH model makes estimates more precise, although it works much better with incomes than with satisfaction.

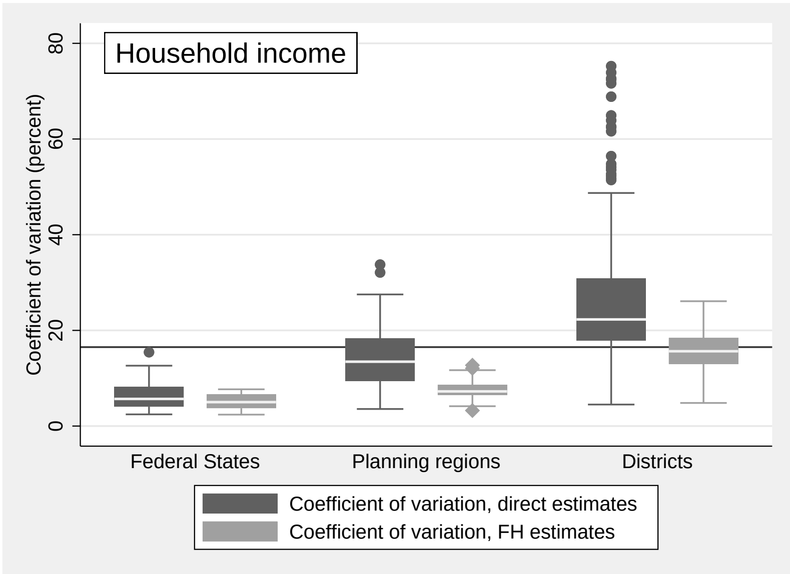

Note: The horizontal line indicates the precision threshold of 16.5 percent. Only in-sample domains are plotted. Eight domains with coefficients of variation larger than 80 percent are not shown. Data are from SOEP v34. Own computations.

Figure 3: Boxplots for the distribution of the coefficients of variation for equivalent HH pre-tax income for the federal states, the planning regions, and the districts.

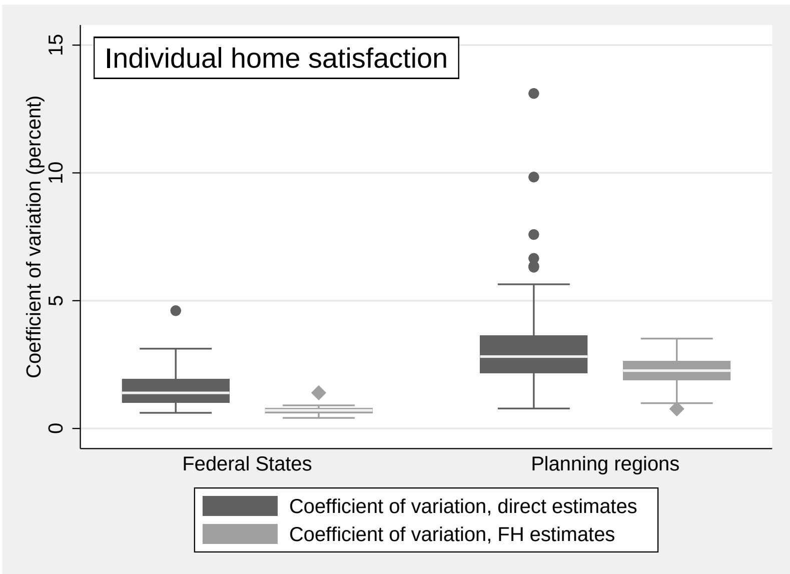

Note: Data are from SOEP v34. Own computations.

Figure 4: Boxplots for the distribution of the coefficients of variation for individual satisfaction with home for the federal states and planning regions.

3 Conclusion

The two SOEP-based applications show how Stata's `fayherriot` package works and its benefits. The package is easy-to-use and offers the opportunity to improve the precision of domain-level indicators by adding auxiliary data from registers or administrations. In particular, if the direct estimates suffer from big variation and the residuals follow a normal distribution, as is the case for income variables, the FH estimator leads to far more precise results than direct estimation. Nevertheless, the smaller the regional level, the lower is the precision, even if using `fayherriot`. In contrast, if the residuals are not normally distributed, the assumptions of the FH model are not fulfilled and a suitable variable transformation or other models should be considered.

Thus, the new Stata package `fayherriot` is an appropriate tool for minimizing problems with low precision of estimates at the domain-level, although its functionality depends on the structure of the variable of interest.

References

Bug, M., M. Kroh, and K. Meier (2015). Regional crime rates and fear of crime: Wisind findings. *DIW Economic Bulletin 5*(12), 167–176.

Bundesinstitut für Bau-, Stadt-, und Raumforschung Bonn (2018). Indikatoren und Karten zur Raum- und Stadtentwicklung. `https://www.inkar.de/`. Accessed: 27.05.2019.

Deckers, T., A. Falk, and H. Schildberg-Hörisch (2016). Nominal or real? the impact of regional price levels on satisfaction with life. *The BE Journal of Economic Analysis & Policy 16*(3), 1337–1358.

Eibich, P. and N. R. Ziebarth (2014). Analyzing regional variation in health care utilization using (rich) household microdata. *Health Policy 114*(1), 41–53.

Eilers, H. L. (2019). *Housing prices and neighborhood structure.* Doctoral thesis, Ruhr-Universität Bochum, Universitätsbibliothek.

Eurostat (2013). Handbook on precision requirements and variance estimation for ess households survey. Methodologies and working papers, European Union.

Fabrizi, E. and C. Trivisano (2016). Small area estimation of the Gini concentration coefficient. *Computational Statistics and Data Analysis 99*, 223–234.

Fay, R. E. and R. A. Herriot (1979). Estimates of income for small places: an application of james-stein procedures to census data. *Journal of the American Statistical Association 74*(366a), 269–277.

Goebel, J., M. M. Grabka, S. Liebig, M. Kroh, D. Richter, C. Schröder, and J. Schupp (2019). The german socio-economic panel (soep). *Journal of Economics and Statistics 239*(2), 345–360.

Halbmeier, C., A.-K. Kreutzmann, T. Schmid, and C. Schröder (2019). The fayherriot command for estimating small-area indicators. *The Stata Journal 19*(3), 626–644.

Kreutzmann, A.-K., S. Pannier, N. Rojas-Perilla, T. Schmid, M. Templ, and N. Tzavidis (2019). *emdi: Estimating and Mapping Disaggregated Indicators.* R package version 1.1.5.

Krieg, O. and B. Raffelhüschen (2017). *Deutsche Post Glücksatlas 2017.* Knaus Verlag.

Molina, I. and Y. Marhuenda (2015). sae: An R Package for Small Area Estimation. *The R Journal 7*(1), 81–98.

Piacentini, M. (2014). Measuring income inequality and poverty at the regional level in OECD countries. *OECD Statistics Working Papers*. No. 2014/03.

Rendtel, U. (1995). *Lebenslagen im Wandel: Panelausfälle und Panelrepräsentativität*, Volume 8. Campus Verlag.

Rojas-Perilla, N., S. Pannier, T. Schmid, and N. Tzavidis (2019). Data-driven transformations in small area estimation. *Journal of the Royal Statistical Society: Series A (Statistics in Society) forthcoming.*

Statistics Canada (2013). 2013 National Graduate Survey (Class of 2009-2010). Microdata user guide. *Statistics Canada*.

Sugasawa, S. and T. Kubokawa (2015). Parametric Transformed Fay-Herriot Model for Small Area Estimation. *Journal of Multivariate Analysis 139*, 295–311.

Tzavidis, N., L.-C. Zhang, A. Luna, T. Schmid, and N. Rojas-Perilla (2018). From start to finish: a framework for the production of small area official statistics. *Journal of the Royal Statistical Society: Series A (Statistics in Society) 181*(4), 927–979.