"An Index of inter-industry wage inequality" by Nathalie Scholl: A Reply*

James Galbraith† Jaehee Choi† Amin Shams‡

University of Texas Inequality Project

Working Paper #74

October 26, 2017

Abstract

Using methods similar to the ones developed by the University of Texas Inequality Project (UTIP), Scholl (2017) has produced a data set that measures inter-industry inequality of wages. In this paper, we review her analysis and show that a major source of differences between her measures and those of UTIP arises from the treatment of the raw UNIDO Industrial Statistics data. First, UTIP adjusts industrial categories to avoid distortions arising from country-specific reclassification of industries in many countries between 1969 and 2015. Second, we pay special attention to coding errors and other data anomalies at the level of industry payrolls or employment, which can have dramatic effects on the between-groups component of Theil's T statistic. Correcting for these small problems produces significantly lower signal-to-noise ratios in UTIP measures compared to Scholl's. Moreover, we find that the UTIP-UNIDO series does a good job of capturing the dynamics of other available income inequality measures, hence the claim that there is no relationship between inter-industry inequalities and household income inequalities is contradicted by our evidence.

JEL Classification: D30, J30, J31, O15
Keywords: Theil statistics, Gini coefficient, income inequality, manufacturing pay inequality

*The authors thank Nathalie Scholl for sharing her measures of inter-industry inequality of manufacturing wages, and for cordial correspondence on these issues over some months.

†LBJ School of Public Affairs, University of Texas at Austin. Email:galbraith@mail.utexas.edu

‡McCombs School of Business, UT-Austin

1 Introduction

Nathalie Scholl (2017) has produced a data set of her own to measure inter-industry inequality of wages – actually, of average payroll per employee. She uses methods similar to ours, as developed by the University of Texas Inequality Project over the years, and specifically as published as the UTIP-UNIDO data set of between-industry pay inequality measures. She has concluded from her comparisons that her data set is not a good instrument for broader measures of inequality, such as measures of income inequality by households. This is all very well. But it supports no inference about the quality of our data set or whether our data set is a useful instrument for broader inequality measures.

Our purpose in this reply is to distinguish Scholl‘s work from our own. This distinction may be muddied for some readers by this sentence on page 2 of her paper: "The estimates suggest that the association between the Theil index of manufacturing and income inequality is neither very stable, nor strong enough to postulate an economically meaningful link between the two concepts," By "*the estimates*" Scholl is referring to the relationship between *her* index and a set of broader measures of income inequality. By "*the Theil index*" she is referring to *her* calculations, not to ours, although ours is also a Theil index of inequality in manufacturing pay. By using the word "*concepts*", she makes an extrapolation from her particular index to the general relationship between a between-industries measure of inequality in average payrolls and broader measures of inequality. This extrapolation is unjustified. The weak (or indeed, absent) relationship is a feature of her particular index alone.

In this paper, we review her analysis and show that a major source of differences between her measures and those of UTIP arises from the treatment of the raw UNIDO Industrial Statistics data. In the next section, we describe this process in greater detail. As in any empirical work, any extreme changes in the raw data deserve special attention. In the third section, we discuss an issue raised by Scholl, concerning the measurement of inequalities within sectors. We then discuss a final issue, namely the relationship between our measures and more general measures of income inequality, in the fourth section. An appendix provides further detail on the data adjustments we made in the course of calculating our between-groups inequality measures from the raw UNIDO Industrial Statistics, and a country-by-country graphical comparison of our measures with those of Scholl, for those countries where there are significant discrepancies between her measurements and our own.

2 Main sources of difference between Scholl and UTIP measures

A comparison of Scholl's measures with our own shows that in the cases where her measures differ significantly from ours, the differences can often be traced to highly erratic and *per se* improbable movements of her measures. Examples include Australia, Austria, Belgium, Botswana, Bulgaria, Estonia, Germany, Haiti, Hong Kong, Jamaica, Jordan, Kenya, Luxembourg, the Netherlands, New Zealand, Norway, Uganda, the UK, the USA and Yemen, as shown in Figures A10-A38 in the appendix to this reply.

These problems in Scholl's data arise from quality issues in the published versions of the UNIDO Industrial Statistics. Although that data set is a valuable and generally reliable compilation of consistent industrial data supplied by many member states to UNIDO under a common accounting scheme, it suffers from occasional missing entries in the underlying reporting of employment or payrolls, and from some apparent recording errors, such as extra digits and misplaced decimal points. The between-groups component of Theil's T statistic is extremely sensitive to these small problems, which show up as spikes, jumps or plateaus in the resulting inequality series, far out of line with the normal historical development. The likely instances of these problems are easy to spot, although locating typos in the mass of underlying payroll/employment data can be a challenge.

We worked carefully to inspect the source data and to clean apparent recording errors and minor omissions from our measures. Our approach to omissions is quite simple, usually a matter of taking an average across a missing year or two, for either payroll or employment (or both). Errors *introduced* by this and similar adjustments are likely to be small, since manufacturing is a continuous activity that rarely (if ever) drops to zero in a year and rebounds in the next; nor do payrolls jump ten-fold and then fall back. Correcting for these small problems generates a characteristically smooth(er) historical series, where major changes are often associated with clear-cut historical events (such as wars, revolutions, economic crises). This appears to be a major source of difference between the two data sets. Evidence for this difference is a lower level of the coefficients of variation through time in UTIP series compared with Scholl's, and the fact that the UTIP series show no statistical difference in those coefficients, as between series which are highly correlated with Scholl's and those that are not. In Scholl's data, due to the erratic behavior of some of her series, these coefficients are significantly higher for series that exhibit low correlation with the UTIP series. Table A1 in the appendix gives the details.

There is also a problem of changes in the category schemes, which occur from time to time as new sectors (such as information technologies) emerge and are accounted for separately in industrial data. We approached the problem of changing category schemes by careful recombination of sectors (reuniting "mother" and "daughter" sectors) so as to preserve consistency over time. Scholl generally favors either curtailing the data series or dropping the new sectors, both of which entail information loss and also distortion, reducing the comparability of measures across countries and through time. In the appendix, we provide examples of how recombining mother and daughter sectors preserves historical continuity in the data and provide the rationale for our procedure (see Figures A1-A9 and Table A1). Scholl correctly states that these differences in approach to categorization are not the main thing that sets her data set apart from ours. The main thing is the erratic behavior of her series, as compared with ours, in the specific cases where the correlation between the two series is low.

In other cases, discrepancies are due to short time series in the Scholl measures, relative to ours; such are the cases of Portugal and Eritrea, for instance. Short coverage in Scholl's data is due to the fact that while we have been calculating these series for years, including from older versions of the UNIDO Industrial Statistics, Scholl had access only to the most recent release. The UNIDO data are not revised over time, so there is usually no reason to recalculate older inequality measures when the data set is updated. Nor is there any reason to discard older measures, merely because the underlying data may not be included in a new release.

3 Inequality "Within-Sectors": A Red Herring

Scholl makes an effort to show the importance of inequality "within sectors." Based on her calculations, she argues that inequality within a sector often moves in the opposite direction from inequality measured between sectors, thus obviating the usefulness of the latter as an instrument for broader inequality measures. To be clear, her measure of inequality within a given (say, 3-digit) sector is not a measure of individual pay in sector; it is, rather, also a measure of inequality *between-sectors*, but measured at the 4-digit level within the 3- digit classification.

There are two distinct issues here. One is whether the level of inequality reported in a Theil measure is affected by the degree of disaggregation. The answer is that it is: the more finely divided the group structure, the more inequality one observes. Put another way, inequality measured across any particular group structure is merely a lower-bound for the inequality of the population encompassed by the groups. This is a point everyone working with this sort of data understands; to divide any group into sub-groups is to move inequality from "within-group", where it is not observed, to "between-groups" where it is.

The other issue – the issue under discussion here – is whether the movement of inequality measured over time between a coarsely-divided set of groups is a good estimate of the movement of inequality measured between groups divided more finely. This question is examined in Conceição and Galbraith (2000) and Conceição, Galbraith, and Bradford (2001), who show that this is the case. Indeed with industrial classifications little information is gained by considering lower-level groups, since the dynamics of the Theil index are preserved at various levels of aggregation.

The essence of this issue is the nature of a classification scheme. Industrial classifications in particular are a type of filing system; they are based on *a priori* categories specified for historical, political and economic reasons. If the classification were entirely at random, a lower-level (more disaggregated) grouping would give more detail, but no more information than the higher level (less disaggregated) classification. In fact the situation is generally different: lower-level groupings (say, 3-digit groups inside a 2-digit group, 4-digit groups inside 3-digit groups) are more similar than the parent groups, and so breaking them out adds progressively *less* to our knowledge of the whole distribution. As a practical matter, beyond a certain (and quite early) point in the process of disaggregation, the return does not justify the cost.

We have worked on this question with fine classifications in the United States (Conceição, Galbraith, and Bradford 2001). In that reliable data environment, before the shift from SIC to NAICS codes, we showed that movement of inequality at the 2-digit level was essentially identical to movement calculated from 3- or 4-digit classifications. Scholl finds many cases where finer categorizations yield inequality measures that are uncorrelated, over time, to the inequalities found by the coarser disaggregation. The question, then, is why is this the case? The apparent answer is that in international data the more finely disaggregated data are simply erratic and unreliable, and the problems that we have labored over in the case of relatively-highly-aggregated groups, namely empty cells and coding errors, are multiplied by the larger number of groups involved.

After the SIC-to-NAICS shift, even US data became problematic at lower levels of aggregation, for instance at the 3-digit level and beyond. The reclassification issue aside, there are also many sudden category changes made at the lower level. These changes were often not documented, and we could find out about them only by making repeated inquiries to the Bureau of Economic Analysis. In international data, while in principle it is possible to work with lower level industries to obtain inequality within a 2-digit or 3-digit sector, it would require a tremendous investment of time and resources to go over individual countries at the 3- or 4-digit industry levels, in order to find and treat coding errors and other anomalies. In practice it would be almost impossible to account for any extreme changes in Theil values without close cooperation from the statistics offices of individual countries. Given limited resources, we do not believe there is a pay-off from attempting to refine the international inequality measures in this way.

4 Manufacturing Wage Inequality and Household Income Inequality: How Close Are They?

Finally, the question of comparability between the *concepts* of inequality in average payroll measured across manufacturing and inequality in (say) household incomes must be addressed. We have often encountered reasonable skepticism on this point, since manufacturing employment is often just a small fraction of total employment, and wage income only part of total income. Entire sectors, from finance to farming, and not to mention the informal economy, are absent from our source data. How can a measure of inequalities between coarsely-divided manufacturing sectors act as an effective instrument for the inequality of a proper random sample of households?

The answer is in four parts. First, as discussed above, the movement of inequalities between groups *is* an effective proxy for the movement of inequalities within those same groups; this is the property of self-similarity at different scales, familiar to students of fractal geometry. (It is also the reason why photographs of coarse and fine resolution may nevertheless show the same objects.) Second, manufacturing is not isolated from agriculture, services, or finance; it stands in a particular relationship to those other sectors, and an increase in the inequality within manufacturing will generally reflect an increase in the gap between manufacturing and agriculture, or manufacturing and services, or manufacturing and finance. Third, manufacturing tends to be the active element in the movement of income distributions; if one isolates peasant farming or low-wage services, you find that the inequalities within those sectors are low. It is not a surprise; there are by definition no rich peasants.

The fourth part of the answer is that the correspondence doesn't always work. In some countries – notably the United States – capital incomes loom large in the income data but are not captured at all by manufacturing pay. In others, de-industrialization can flatten manufacturing pay while overall inequality rises. One has to be careful – but this caveat applies to any form of empirical investigation.

Our assessment of the actual match between our Theil measures and broader inequality measures can be presented in two ways. One is a statistical summary, which shows that our data do in fact match reasonably well with broader measures of economic inequality. This summary is presented in Table 1. Using our measures rather than Scholl's, we replicated her Table 4, in which she argues that *her* index does not have any substantive relationships with overall inequality measures from major global data sets such as the World Income Inequality Database (WIID), the Luxembourg Income Study (LIS), and EU Statistics on Income and Living Conditions (EU SILC). Table 1 exactly matches Scholl's Table 4, but with very different results.[1] Our measures do have statistically-significant relationships to the major income inequality datasets, and the coefficient magnitude is ten-fold greater than for Scholl. The strength of statistical significance decreases for LIS, SILC, and OECD samples; these data have smaller time and geographical coverage and thus smaller sample sizes.

To be sure, the magnitude of the effect of a move in the Theil statistic on a Gini coefficient is quite small; for instance, with respect to the WIID, a one percent increase in the Theil index corresponds to a rise in the Gini of just 0.05%. The reasons for this are two-fold: first that the Theil index between large groups is quite volatile, with changes of ten percent or more not unusual, and doublings not unheard-of. Second, the Gini coefficient is by definition bounded between zero and unity, and in practice largely limited to a small part (say, a fifth) of that range for any given country. Thus a doubling of our Theil index would lead to an increase in a Gini of 0.4 to 0.42, but the latter is not an inconsiderable amount.

Table 1: UTIP-UNIDO as compared to other data sets

Replication of Table 4: Relationship between Theil and Gini

	(1) WIID	(2) SWIID: Net	(3) SWIID: Market	(4) LIS	(5) SILC	(6) WDI	(7) OECD
ln(Theil)	0.049***	0.067***	0.064***	0.038*	0.034*	0.038***	0.026*
	(4.40)	(14.97)	(13.48)	(1.92)	(1.66)	(3.51)	(1.73)
country f.e.	yes	yes	yes	yes	yes	yes	yes
year f.e.	yes	yes	yes	yes	yes	yes	yes
Observations	718	2,092	2,092	173	241	548	232
R^2	0.286	0.258	0.294	0.572	0.156	0.185	0.506
# Countries	60	122	122	36	29	95	232

t statistics in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

Dependent variable = ln(Gini)

Note: We do not have control variables such as urban population share and WIID dummies in our regression, but this would not substantially change the results.

Scholl expresses a further reservation (2017, p. 34) about the incorporation of cross-

1. For the purposes of this replication we used the 2005 version of our data.

country variance into our estimates of the relationship between Theil and Gini measures. Suffice to say, the test of the estimating equation is whether the resulting estimate for gross household income inequality is consistent with the evidence from other sources. We have examined this proposition directly and in detail, and while there are, to be sure, some imperfections, we believe the record speaks for itself.

To effect this examination we made a detailed, country-by-country comparison of our EHII estimate with the full range of reported inequality measures. We have done this by first using our Theil measure of inter-industry payroll inequalities as the building block – along with the ratio of manufacturing employment to population – in a simple model of gross household income inequality, calibrated to an external data set of sample-based measures and controlling for the basic conceptual categories: households or persons, income or consumption, gross or net of tax. We then used the regression coefficients to generate a large data set of estimated household income inequality measures in Gini-coefficient format, the EHII data set which is in wide use. The EHII numbers can then be compared directly to measures of income inequality from the survey literature.

We have carried out this exercise for over thirty countries, in an extensive survey published by the World Bank (Galbraith et al. 2016b; Galbraith et al. 2016a). This shows that our very simple model to translate manufacturing pay inequality into gross household income inequality corresponds well with available survey measures in many countries. An example – chosen because the correlation between our measure and Scholl's is actually negative – is Germany, shown in Figure 1. In this case, our equation yields a Gini estimate that tracks the available survey literature for household gross income inequality to within a Gini point, give or take, over thirty years. Market income inequality measures are uniformly higher, and disposable income inequality measures are uniformly lower than the measurements of gross income inequality; however the trends in all three data types are broadly similar.

Figure 1: Diverse measures of income inequality for Germany

Note: The black line is the EHII estimate from UTIP for gross income inequality. The blue line is the one available direct measure of gross household income inequality for Germany. The dotted black line is the EHII estimate for East Germany (DDR).

There are no perfect measures of inequality, and our data set, like all data sets in this area, is a work in progress. But a conclusion drawn by Scholl from her data set to the general applicability of the Theil method applied to manufacturing data is plainly invalid. It is quite easy to check whether one can calculate a reasonable estimate of household income inequality from inequalities across sectors in manufacturing pay, and it turns out that the answer is, yes you can.

References

Conceição, Pedro, and James K Galbraith. 2000. Constructing Long and Dense Time-series of Inequality Using the Theil Index. *Eastern Economic Journal* 26 (1): 61–74.

Conceição, Pedro, James K Galbraith, and Peter Bradford. 2001. The Theil Index in Sequences of Nested and Hierarchic Grouping Structures: Implications for the Measurement of Inequality through Time, with Data Aggregated at Different Levels of Industrial Classification. *Eastern Economic Journal* 27 (4): 491–514.

Galbraith, James K, Jaehee Choi, Béatrice Halbach, Aleksandra Malinowska, and Wenjie Zhang. 2016a. A Comparison of Major World Inequality Data Sets: LIS, OECD, EU-SILC, WDI, and EHII. In *Income inequality around the world*, 1–48. Emerald Group Publishing Limited.

Galbraith, James K, Béatrice Halbach, Aleksandra Malinowska, Amin Shams, and Wenjie Zhang. 2016b. The UTIP Global Inequality Data Sets 1963-2008: Updates, Revisions and Quality Checks. In *Inequality and growth*, ed. Ove Grandstrand, 7–39. Also appearing as: The UTIP Global Inequality Data Sets, 19632008: Updates, Revisions and Quality Checks in Joseph E. Stiglitz and Martin Guzman, eds., Contemporary Issues in Microeconomics, Basingstoke: Palgrave Macmillan, 2016, pp. 739. Basingstoke: Palgrave MacMillan.

Scholl, Nathalie. 2017. *An Index of Inter-industry Wage Inequality: Trends, Comparisons, and Robustness.* Technical report. Courant Research Centre: Poverty, Equity and Growth-Discussion Papers.

UNIDO. *United Nations International Development Organization Industrial Statistics Database.* Various years.

Appendix: Detailed Notes on Updating the UTIP-UNIDO Data Set

We updated the UTIP-UNIDO data based on the 2012 version of the INDSTAT2 database as well as the 2005 version of the UTIP-UNIDO data. We have made two types of adjustments. One of them involves re-combining sectors that become separated by the introduction of new industrial categories. The other involves the treatment of missing observations and coding errors.

The treatment of missing sectors has a limited effect on the calculated Theil index, as Scholl (2017) confirms. However, among different options, we do a more careful treatment of the missing sectors based on the information reported by the UNIDO on sectors that are combined over time. In this way, we exploit all the information available and do not miss any data point, as opposed to Scholl, who drops either sectors or country-year observations when facing a missing sector. Our time-series data is still consistent within each country because the number of sectors for each country does not change over time. Whether the resulting Theil measures are comparable across countries is another question, especially when there are some differences in the number of categories used for different countries. The correspondence of the derivative EHII measures to sample-based measures of gross household income inequalities suggests that this is not a major issue for advanced or transition countries; it may however contribute to underestimates of inequality in certain developing-country cases.

The second type of adjustment deals with coding errors and other forms of noise in the UNIDO INDSTAT data. Here we compared the calculated data based on the INDSTAT 2012 version (which covers data up to 2008) against the 2005 version of the UTIP-UNIDO data, country-by-country. The data using the 2012 version shows noise and errors even for the data that is already covered in UTIP-UNIDO 2005 version. We replace the new data with the old UTIP-UNIDO data whenever the new data is noisy, taking into account the difference in levels of inequality calculated in the two data sets. Specific examples of such treatment are discussed later. The reason that the data using the INDSTAT 2012 are different from UTIP-UNIDO 2005 can be either of two possibilities. Either, a) The INDSTAT data quality used to be good, but they have changed the data, or reproduced it poorly, in a way that the data points before 2005 were sometimes deteriorated in later releases; or b) that the INDSTAT data were of poor quality beforehand, but were corrected in later releases. There is no evidence that earlier data have been upgraded, and our implicit assumption is that option a) is the case.

A1 Missing Sectors and Re-groupings

The data for certain sectors of the economy in certain countries are only reported in certain years and are missing in other years. For example, in Canada, data for sector #35 (Other transport equipment) start from 1990 and are missing before, while the rest of the sectors start as early as 1963. In order to keep the time-series of Theil index consistent over time, we need to keep the number of industries fixed within each country. So an option could be to drop sector 35 for Canada in the entire sample period. Another option is to drop the country-year observations that have a missing sector, which means dropping observations before 1990 for Canada.

These are the two options considered in Scholl, called long and short versions respectively. However, both of these treatments result in losing information and endangering the accuracy of the data as shown below. An alternative solution used in the UTIP-UNIDO data utilizes the information on sector combinations issued by UNIDO. Certain countries in certain years report the data for a combination of the sectors, which is clearly flagged in the INDSTAT data. For example, it is reported that the data for sector 35 in Canada was combined with sector 34 (Motor vehicles, trailers, semi-trailers) before 1990. The graph below shows the number of employees in sectors 34 and 35 in Canada over time, and clearly confirms such a pattern.

Figure A1

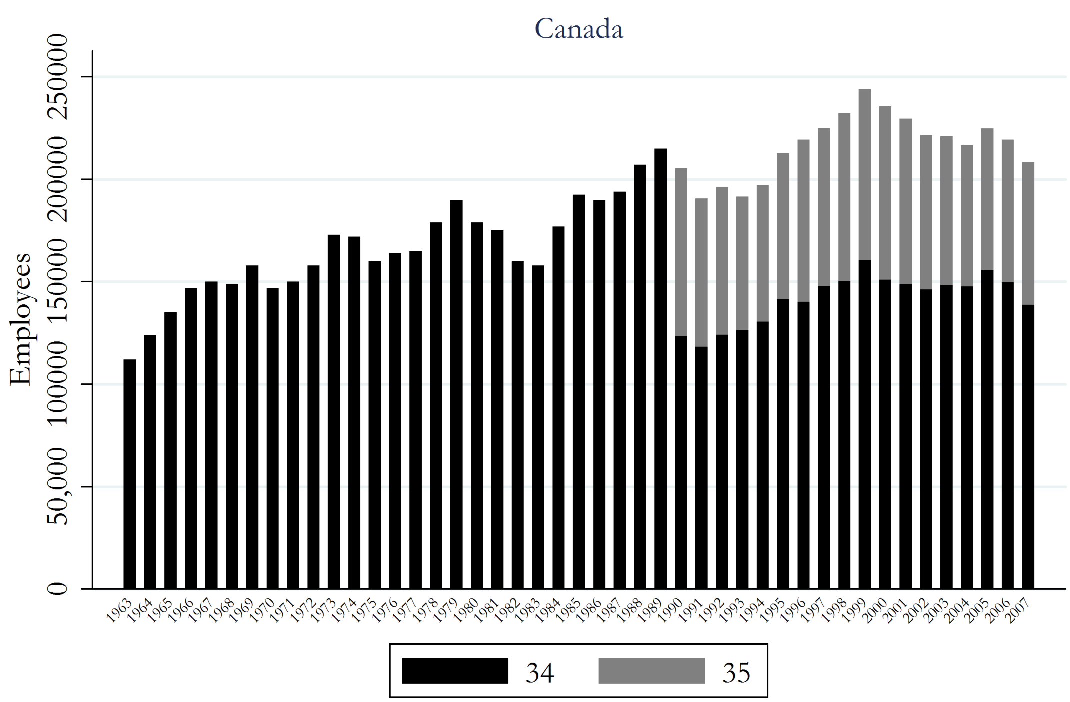

The number of employees reported in sector 34 (the blue bars) drops from more than 200,000 in 1989 to around 120,000 in 1990. The reason is that the data were reported as a combination of sectors 34 and 35 before 1990, but separately after. Simply dropping sector 35, as done in Scholl (2017), suddenly reduces the weight of transportation equipment in the economy significantly and makes the data inconsistent before and after 1990. What we do in UTIP-UNIDO is to aggregate both the number of employees and wages for sectors 34 and 35 for Canada over the entire sample. The graph below shows a similar pattern for the aggregate wages in the two sectors.

Figure A2

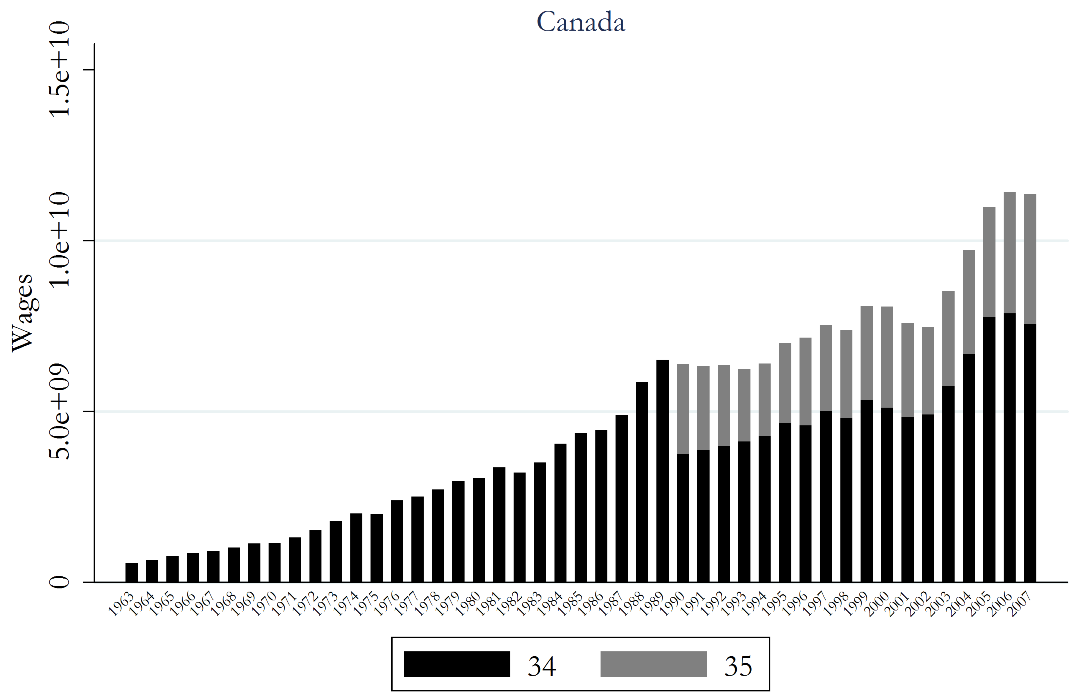

The following graphs show a few other countries with the same pattern:

Figure A3

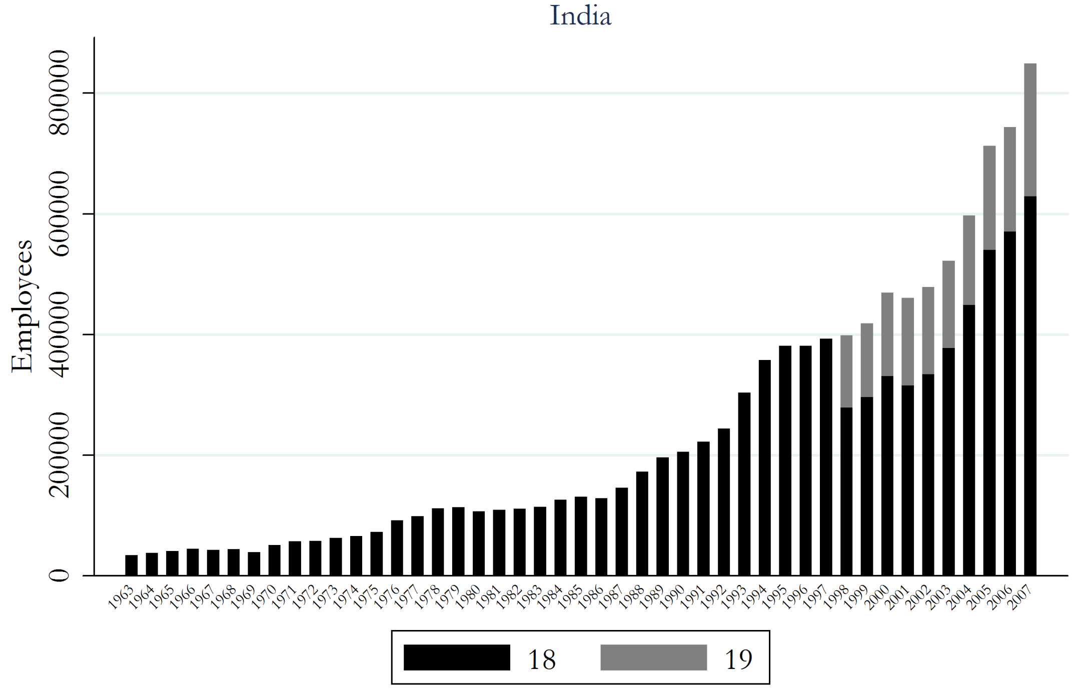

Figure A4

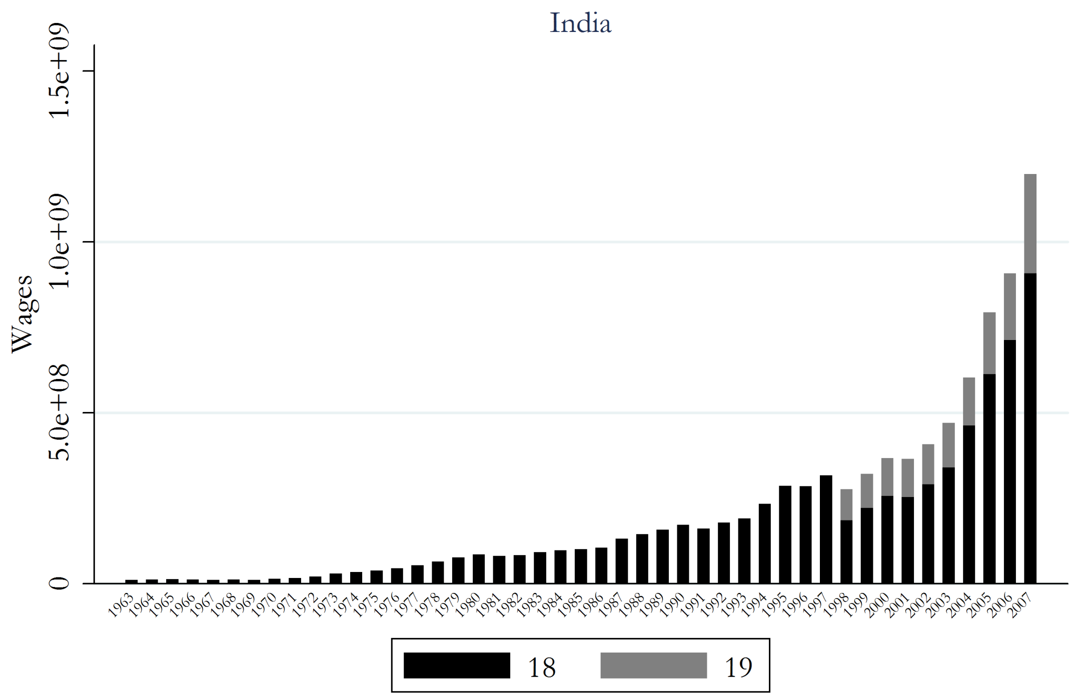

Figure A5

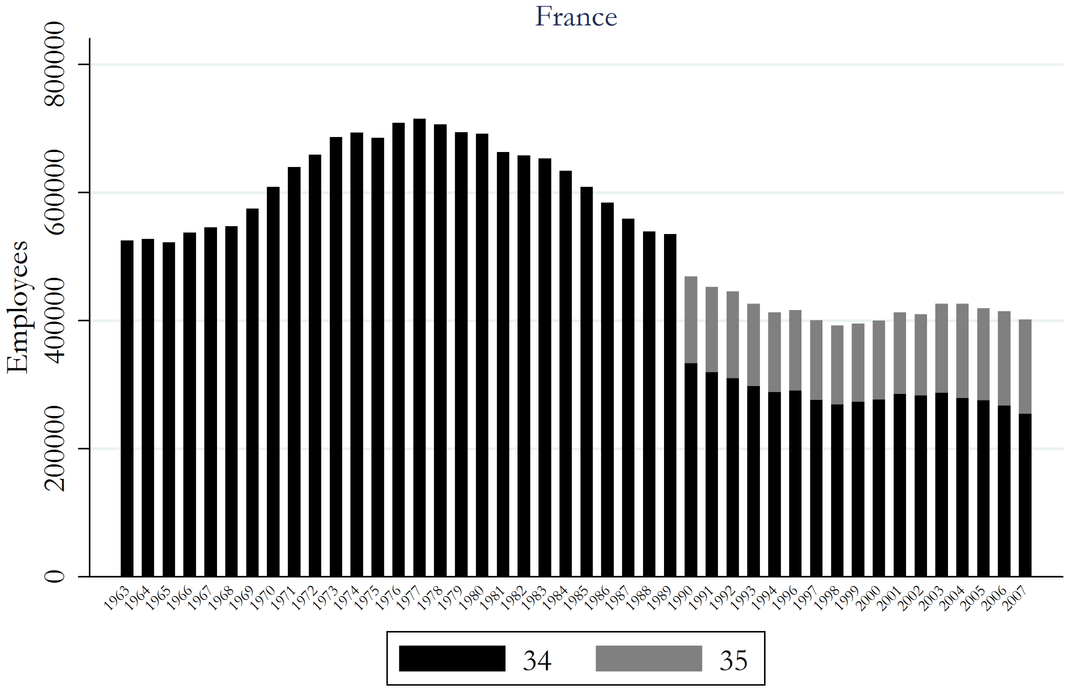

Figure A6

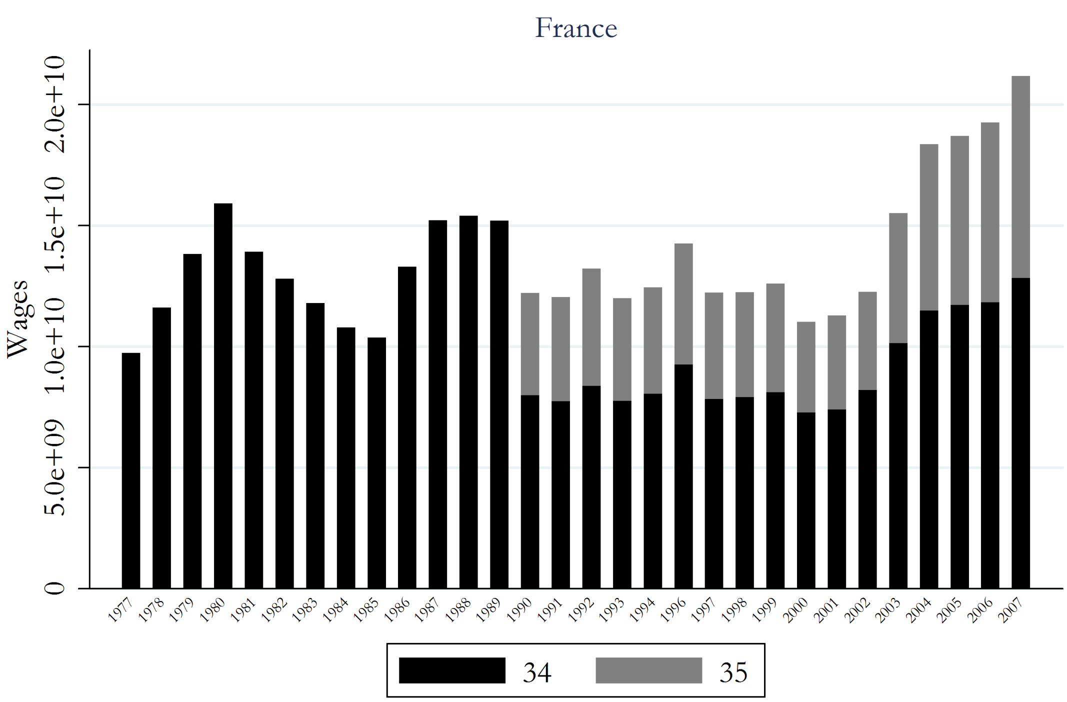

Although the majority of the sector combinations in the data are of a similar pattern, there exist more complicated cases as well. Suppose that sectors A and B are combined before 2000 for a certain country, and sectors B and C are combined at some point after 2000. In a very few cases the number of combined sectors can go up to 8 industries. Such cases mostly happen for sectors 27 to 35. The graphs below shows such examples:

Figure A7

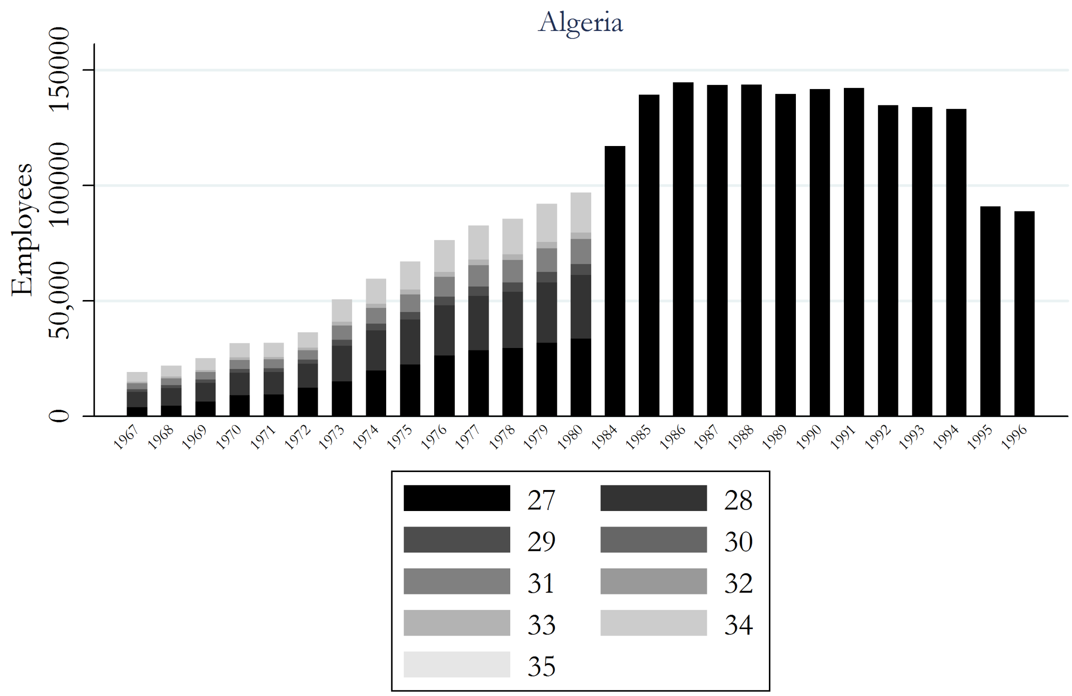

Figure A8

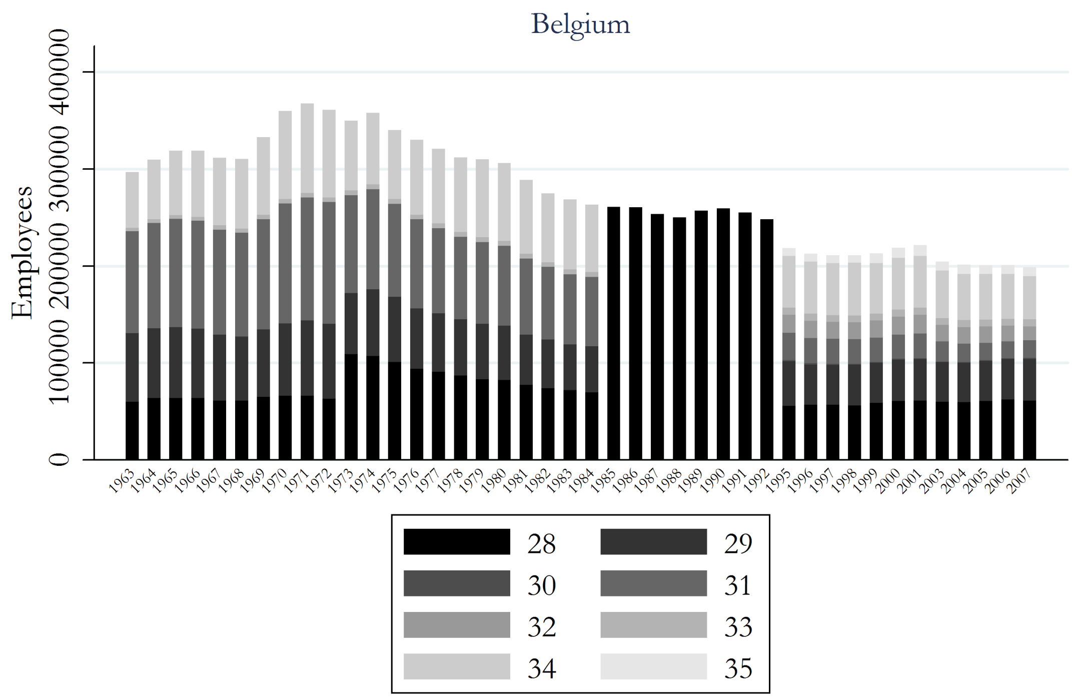

Figure A9

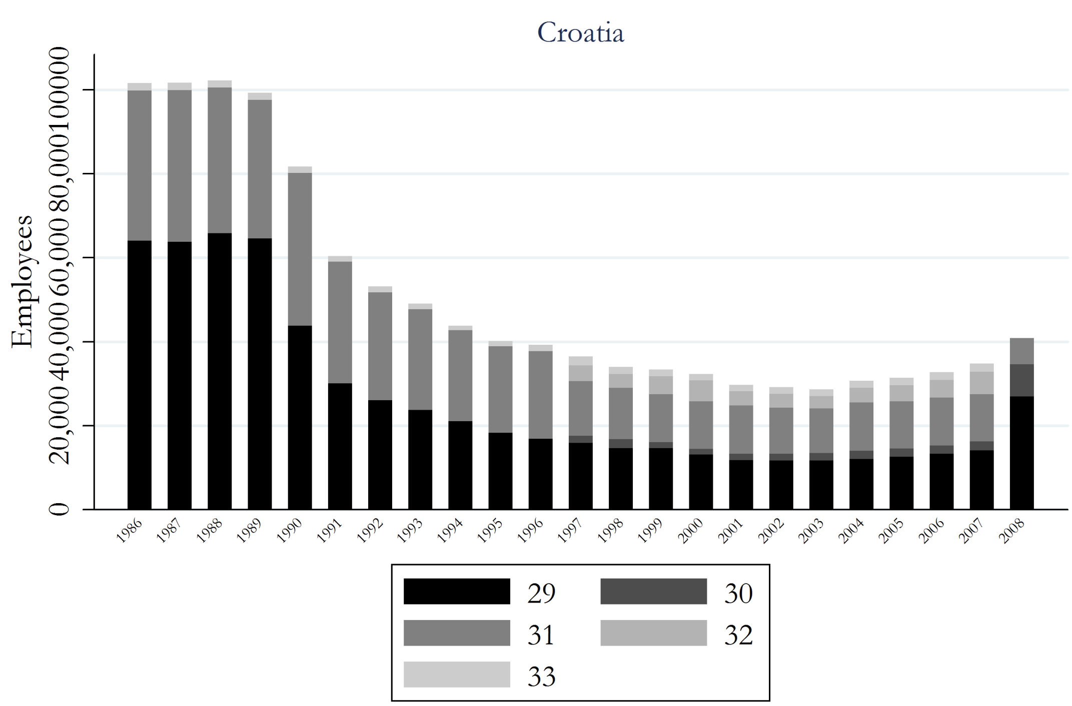

Note that mapping of combined sectors to the reported sectors can vary, depending on how the combination of sectors change in different version of the UNIDO database. But there is a very logical framework behind the combinations that makes the classification uniquely and systematically replicable: whenever two or more sectors are combined in part of the sample, one has to combine their information for the entire sample.

A2 Adjustments for Coding Errors and Other Noise

An accurate treatment of the combined sectors, while crucial, leads to a limited improvement in the data. There are other forms of noise and errors in the data that need treatment. To show how such noise can affect the quality of the estimated Theil index, we start with a comparison of UTIP-UNIDO data with Scholl‘s data. Table 3.A.2 of Scholl reports the mean and standard deviation of her calculated (between-groups component of the) Theil index for each country, as well as the correlation of her index with the UTIP-UNIDO data for each country. Most countries have a high correlation between the two datasets, however, there are also several countries with low or even negative correlations. We will argue here that the low correlation for some of these countries is caused by inadequate treatment of noise and errors in Scholl‘s data.

We divide the countries into those with high correlation between the two datasets, with correlation of higher than 0.9, and those with correlation of less than 0.9. We show below that the low correlation for certain countries is caused by high noise in Scholl‘s data for those countries. First, for each country, we look at the "coefficient of covariation" measure, created by dividing the standard deviation of the Theil index on to the average of the Theil index for that country. The measure shows the precision of the data and reflects the inverse of signal-to-noise ratio. We then test whether this measure is different for low correlation countries than high correlation countries in both datasets. If the coefficient of covariation is higher for low correlation countries than high correlation countries in Scholl‘s data, we can argue that the low correlation is caused by a high noise for those countries in her data.

Table A.1 below, Panel A, shows that countries that have low correlation between UTIP-UNIDO and Scholl have systematically higher coefficient of covariation, meaning that they have lower signal-to-noise ratio. The results are highly statistically and economically significant and are robust to limiting the sample to countries with at least 10 or 20 years of data. For example, the results show that among countries with at least 20 years of data in Scholl‘s sample, coefficient of variation for countries with a low correlation with UTIP data is around 74% while it is 50% for high correlation countries. The difference is statistically significant with a t-stat of -3.19, which is impressive given the small sample size. So the estimations for low correlation countries are systematically different in Scholl's data.

Panel B shows the same exercise using the UTIP-UNIDO data. As opposed to Scholl's data, there is no significant difference between the coefficient of variation for low and high correlation countries in UTIP data. The two panels together suggest that the quality of UTIP-UNIDO data is consistent across different countries while there are countries in Scholl's data with a high level of noise, which also results in a low correlation with the UTIP-UNIDO data. The low correlation between the two data sets for certain countries is caused because the low correlation countries are systematically noisier than high correlation countries in Scholl data, but not in UTIP data.

Table A1

Panel A: Scholl (2017) Data

	LowCorr	HighCorr	Difference	t-stats	N
Coeff. of Variation	68.87	48.90	-19.97**	(-3.23)	133
Coeff. of Variation (>=10Yr)	71.94	49.10	-22.84***	(-3.65)	125
Coeff. of Variation (>=20Yr)	73.73	49.98	-23.76**	(-3.19)	96

Panel B: UTIP-UNIDO Data

	LowCorr	HighCorr	Difference	t-stats	N
Coeff. of Variation	39.15	45.17	6.02	(1.30)	133
Coeff. of Variation (>=10Yr)	39.90	45.93	6.03	(1.25)	125
Coeff. of Variation (>=20Yr)	41.88	45.12	3.24	(0.66)	96

Note: Coefficient of Variation and the Correlation between the Two Datasets. This table shows the results of a t-test of the coefficient of variation between countries with high correlation in UTIP-UNIDO and Scholl data ($correlation > 0.9$) and those with low correlation. Coefficient of variation for each country is calculated as 100 times the standard deviation of the time-series of Theil index divided by the average Theil index for that country. The data on average and standard deviation of the Theil index in Scholl data are from Table 3.A.2.

Below we show examples of noise in Scholl's data. The 30 countries selected are from

Table 3.A.3. and have a correlation less than 0.8 between the Scholl measures and the UTIP-UNIDO measures of industrial pay inequality.

Figure A10

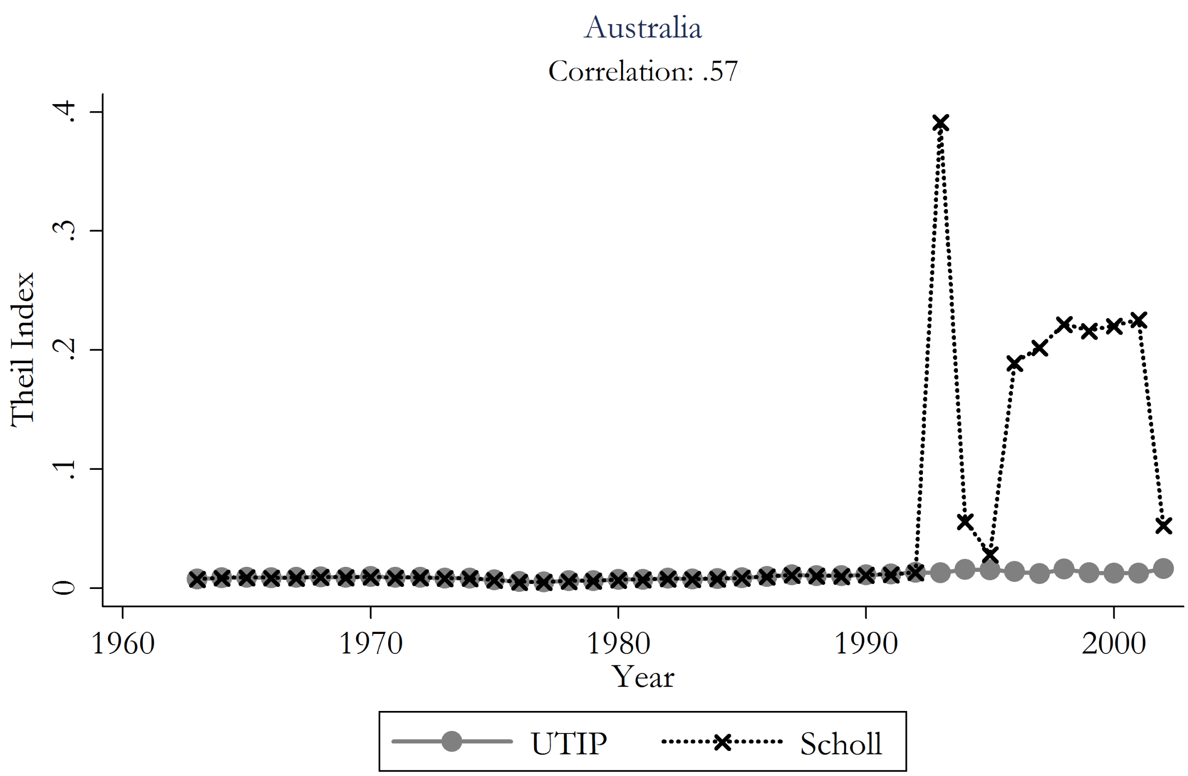

Figure A11

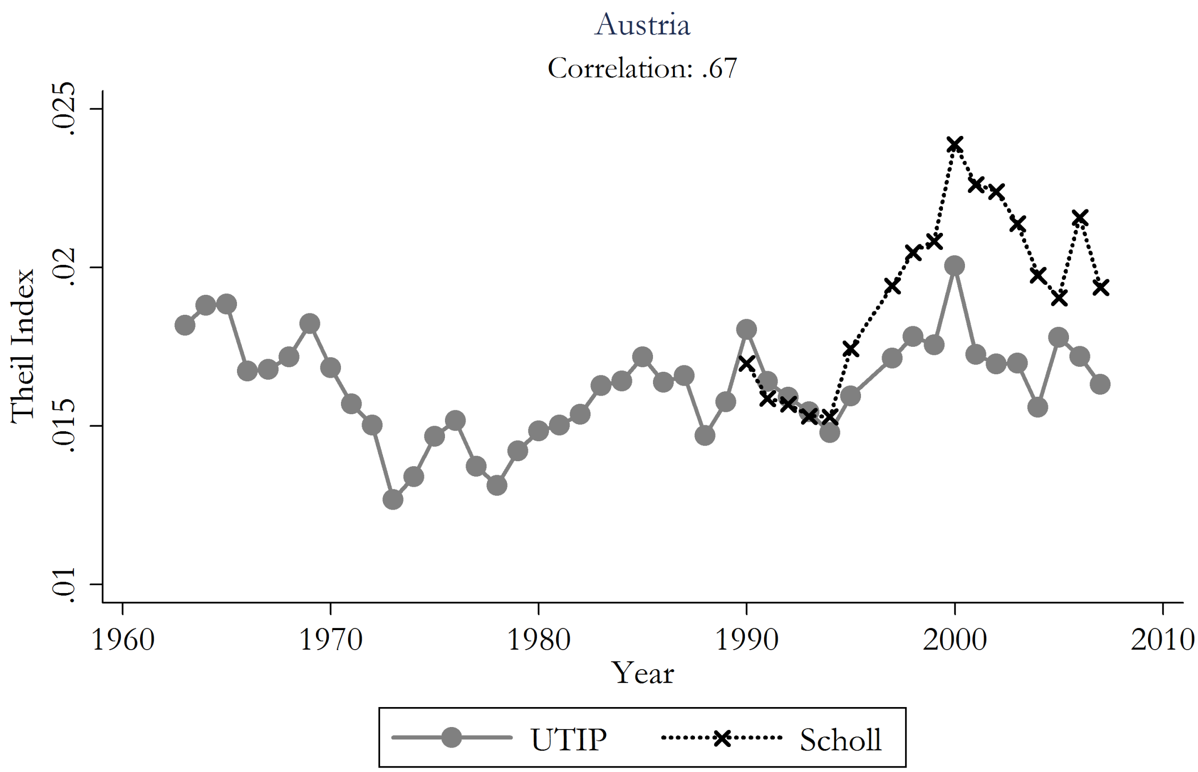

Austria
Correlation: .67
Theil Index
.01
.015
.02
.025
1960
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A12

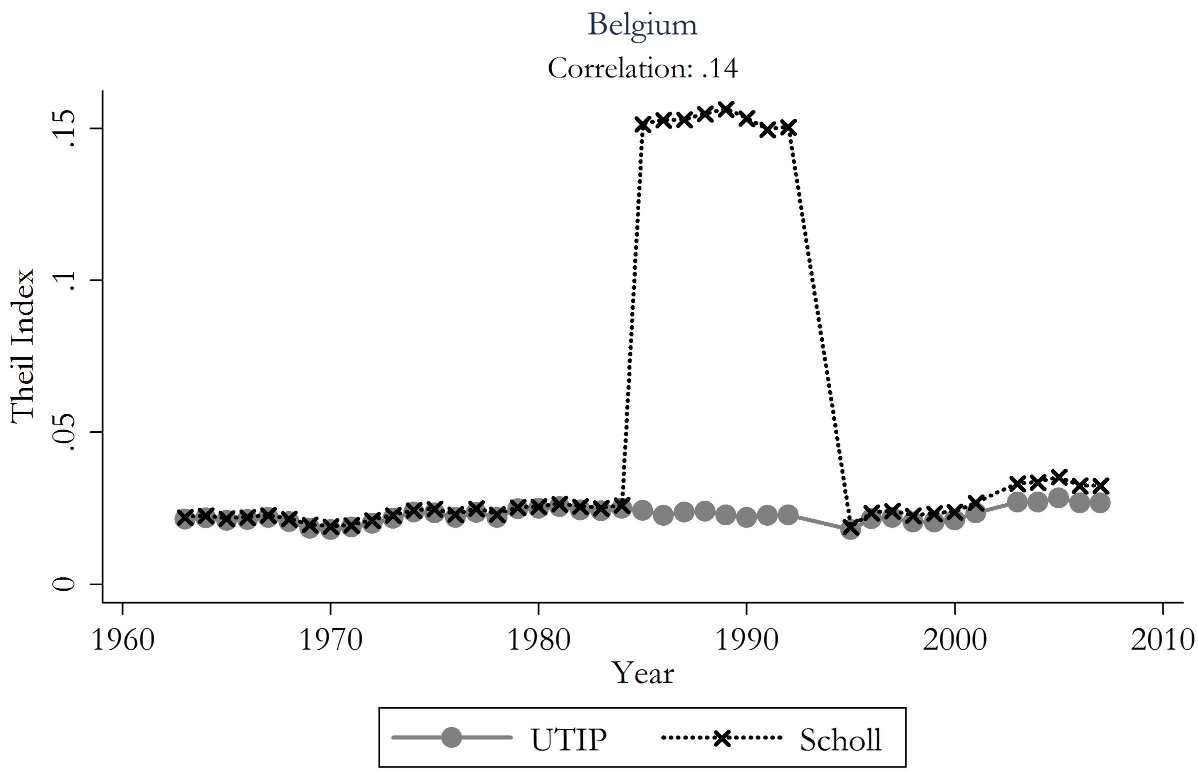
Belgium
Correlation: .14
Theil Index
0
.05
.1
.15
1960
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A13

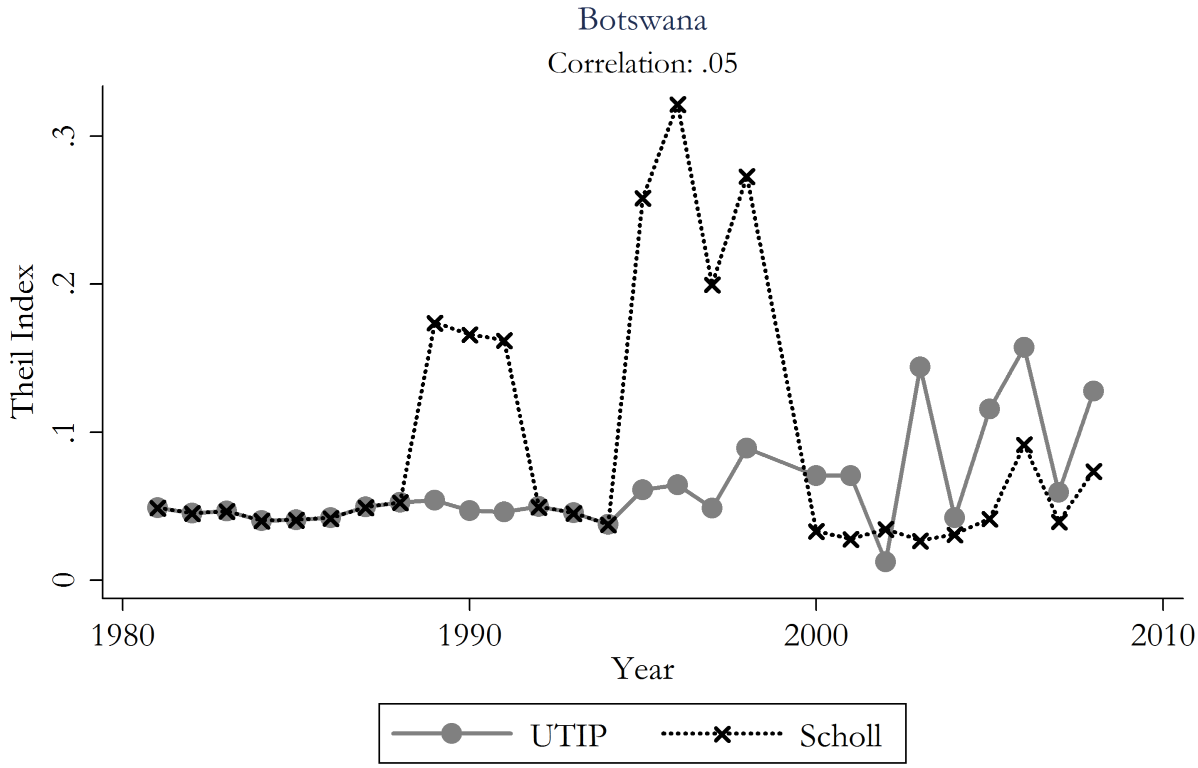
Botswana
Correlation: .05
Theil Index
0
.1
.2
.3
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A14

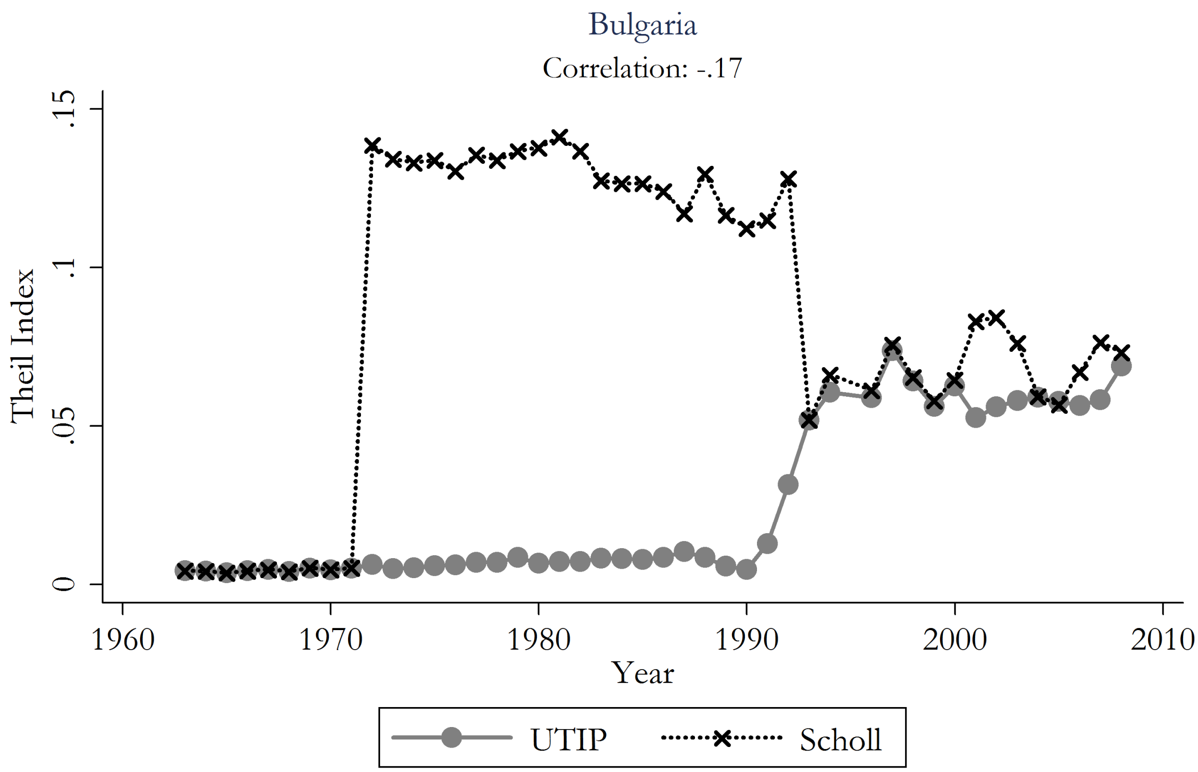
Bulgaria
Correlation: -.17
Theil Index
0
.05
.1
.15
1960
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A15

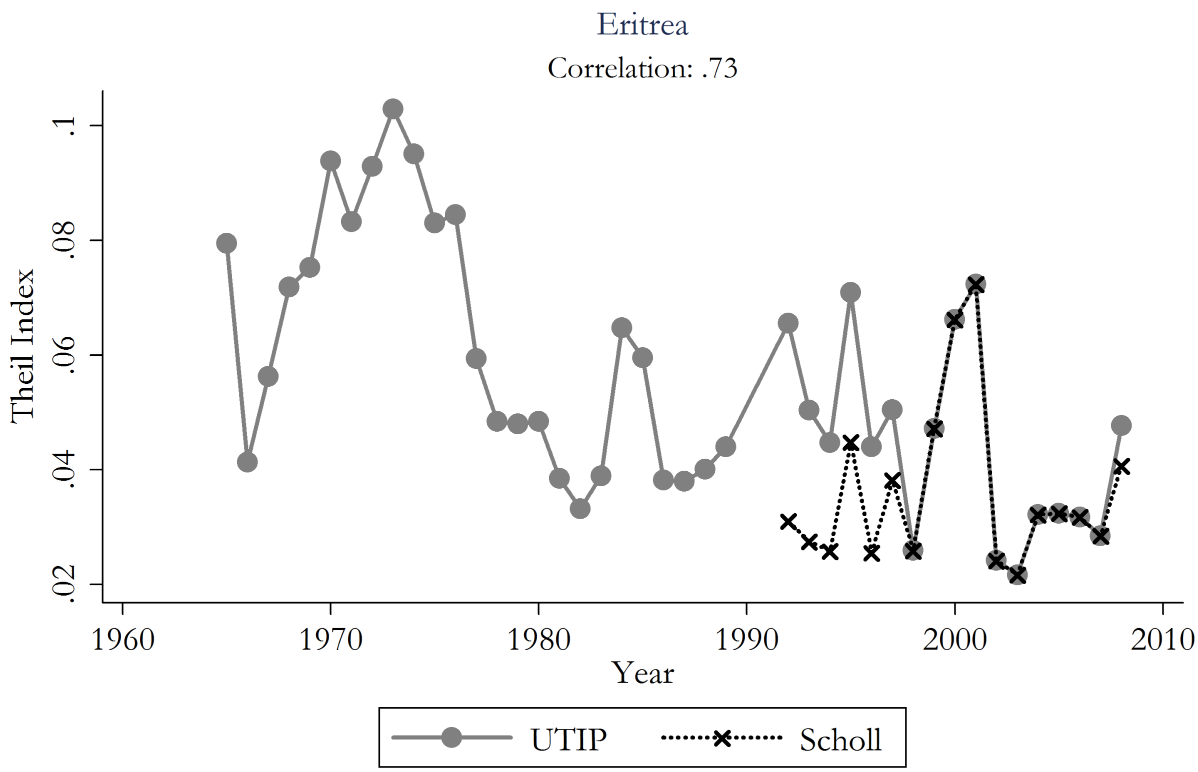
Eritrea
Correlation: .73
Theil Index
.02
.04
.06
.08
.1
1960
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A16

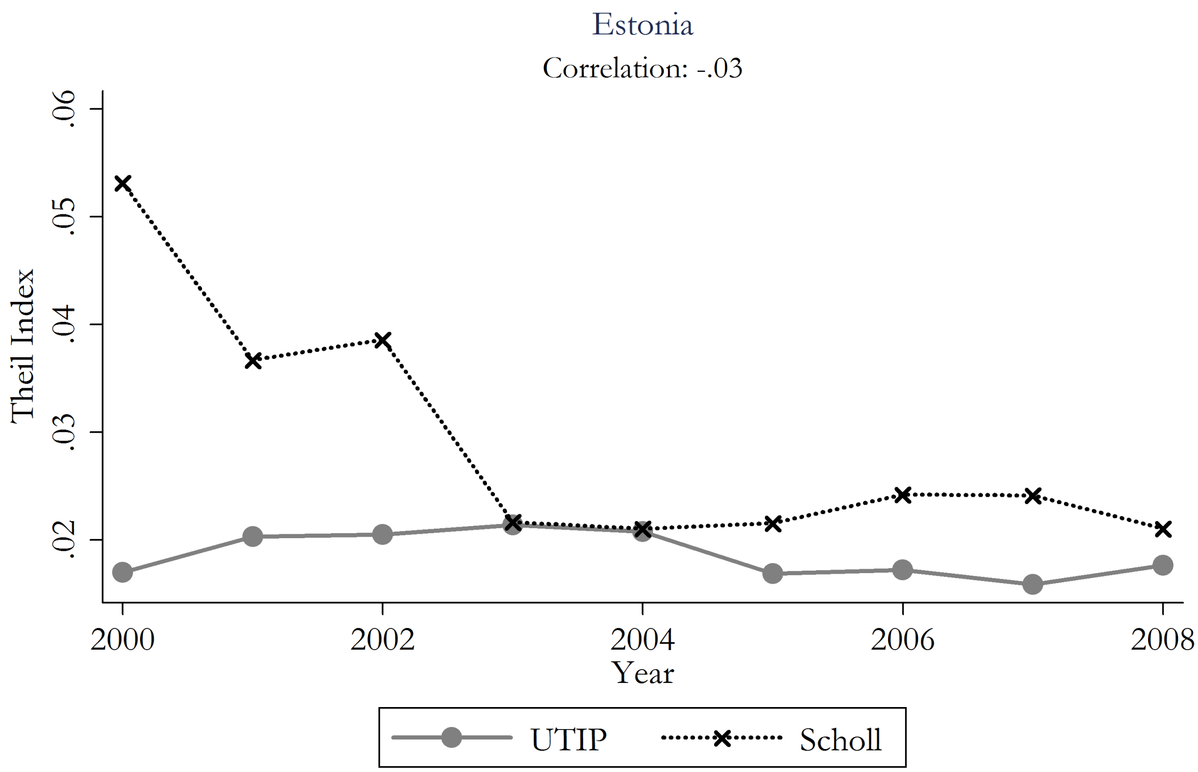
Estonia
Correlation: -.03
Theil Index
.02
.03
.04
.05
.06
2000
2002
2004
2006
2008
Year
UTIP
Scholl

Figure A17

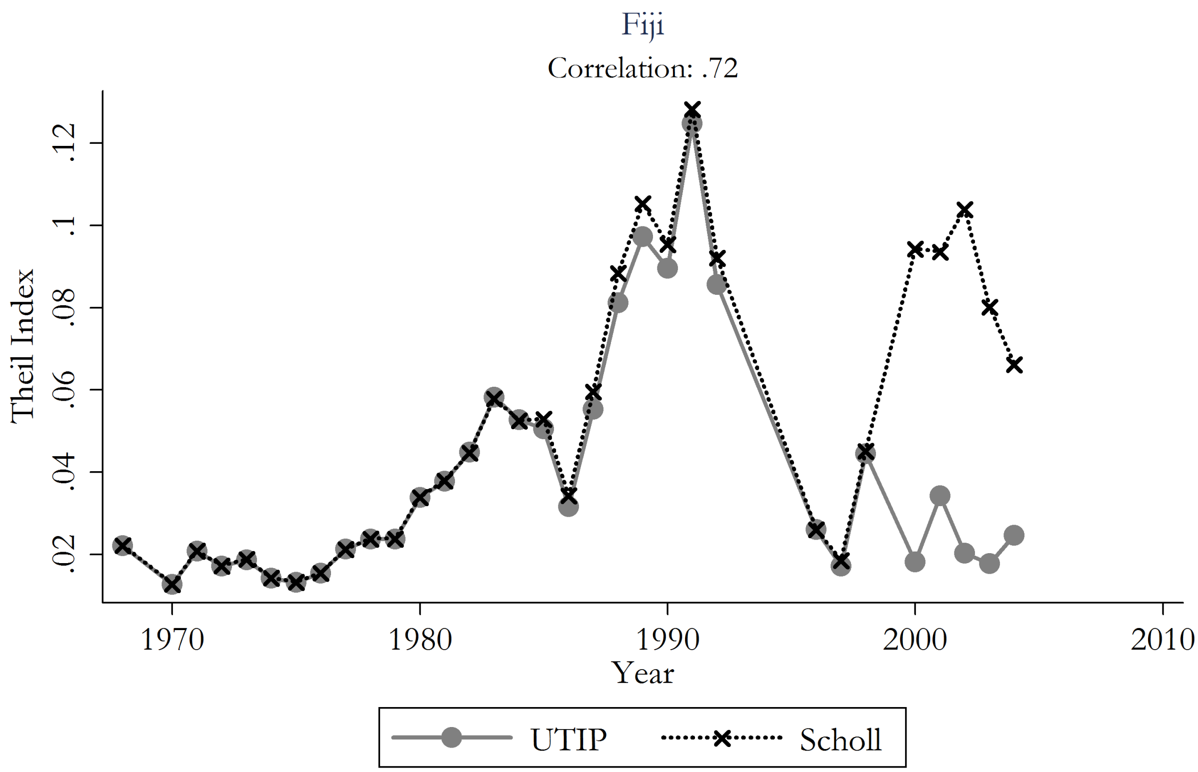
Fiji
Correlation: .72
Theil Index
.02
.04
.06
.08
.1
.12
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A18

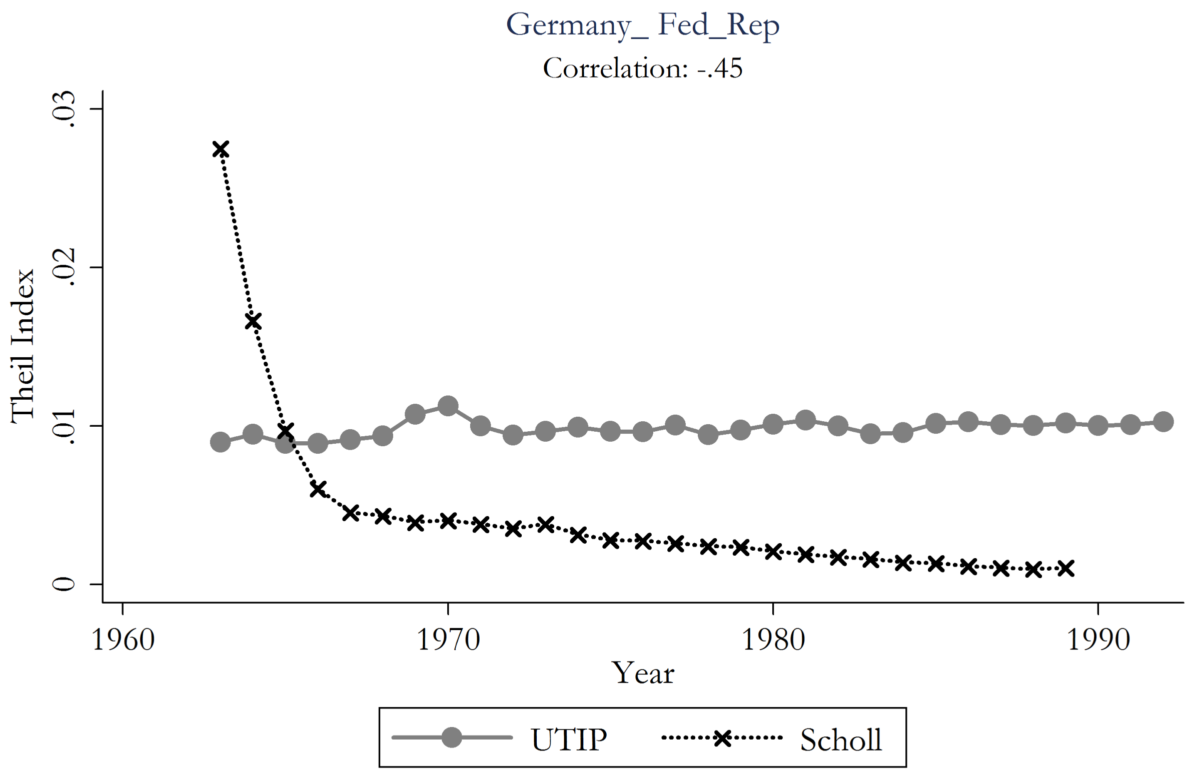

Figure A19

Figure A20

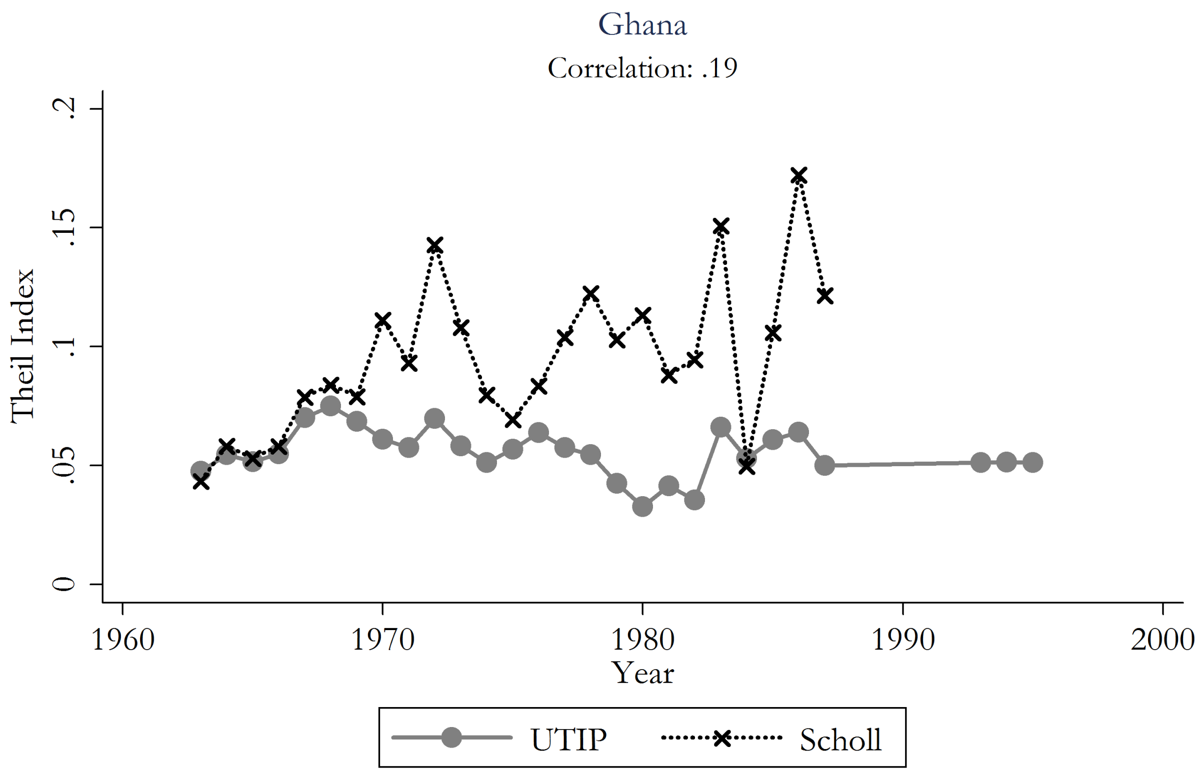

Ghana
Correlation: .19
Theil Index
0
.05
.1
.15
.2
1960
1970
1980
1990
2000
Year
UTIP
Scholl

Figure A21

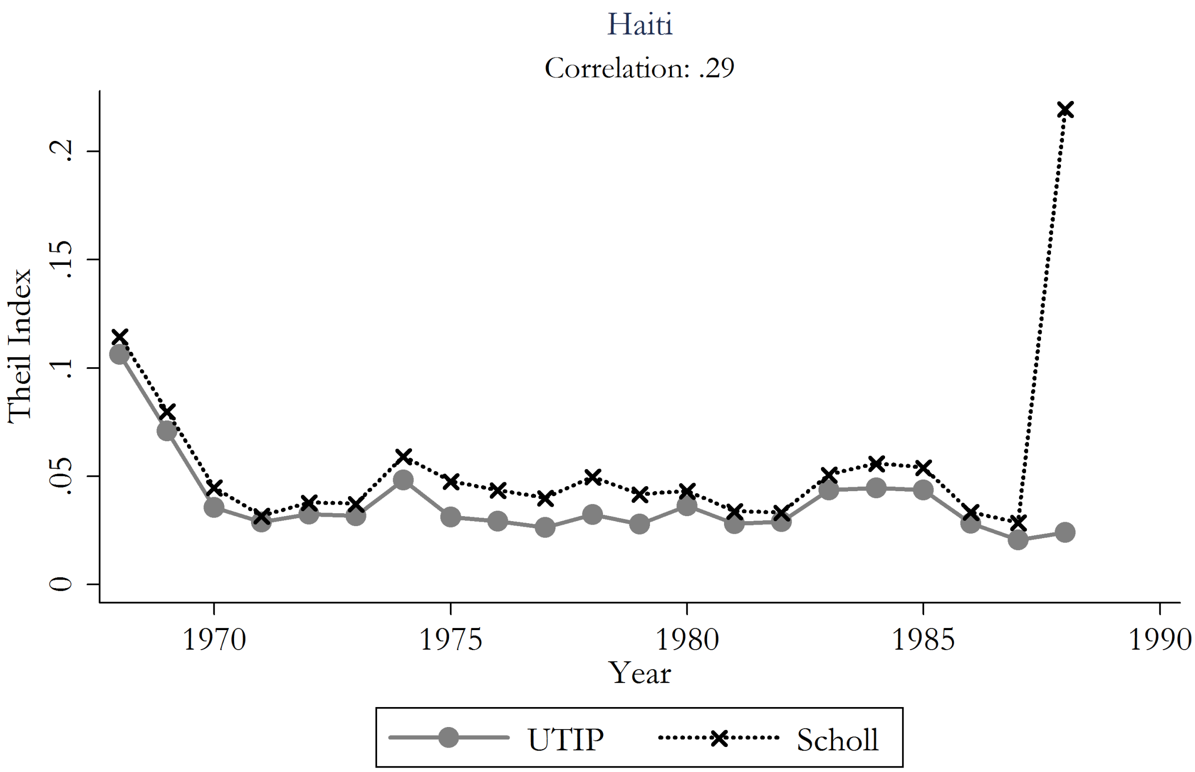
Haiti
Correlation: .29
Theil Index
0
.05
.1
.15
.2
1970
1975
1980
1985
1990
Year
UTIP
Scholl

Figure A22

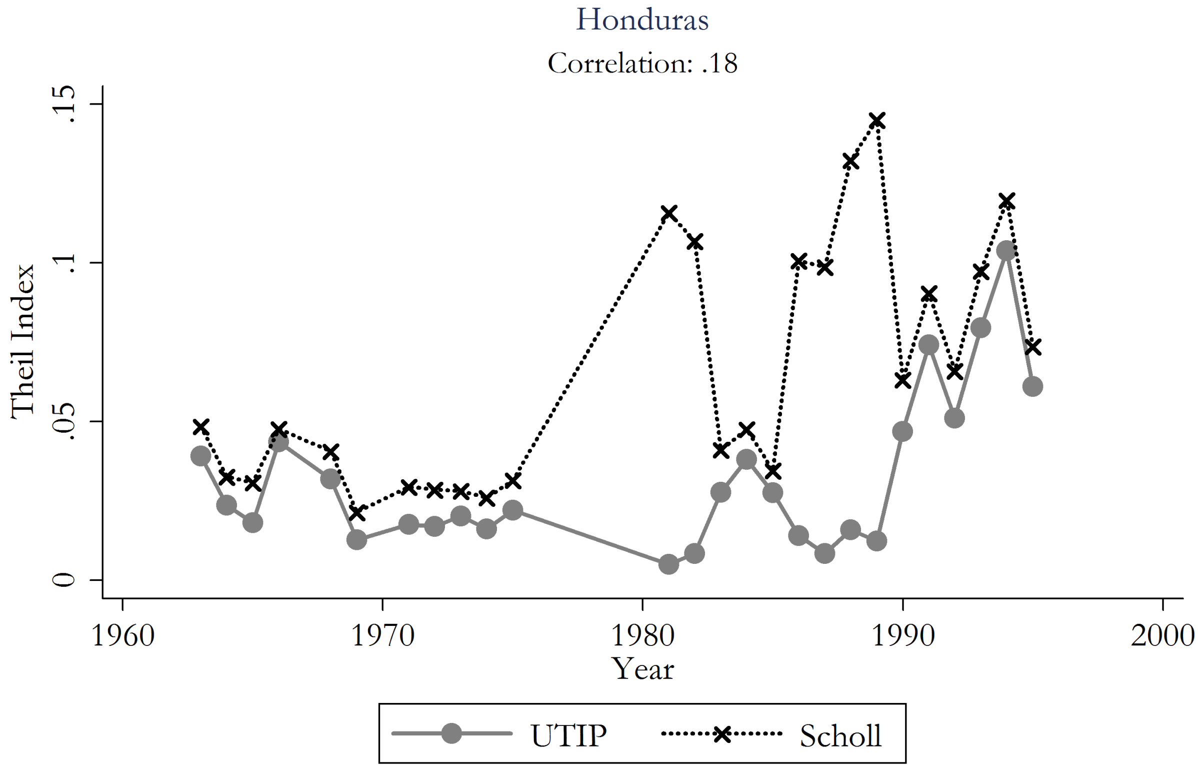
Honduras
Correlation: .18
Theil Index
0
.05
.1
.15
1960
1970
1980
1990
2000
Year
UTIP
Scholl

Figure A23

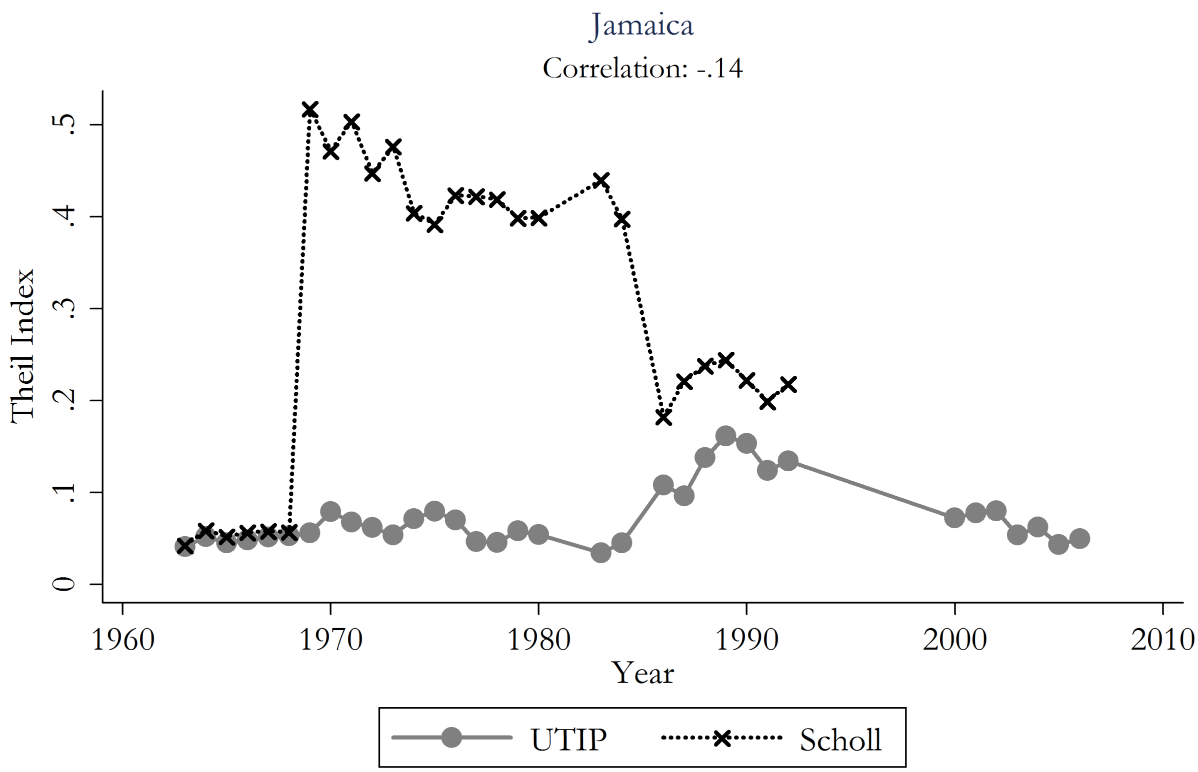

Jamaica
Correlation: -.14
Theil Index
0
.1
.2
.3
.4
.5
1960
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A24

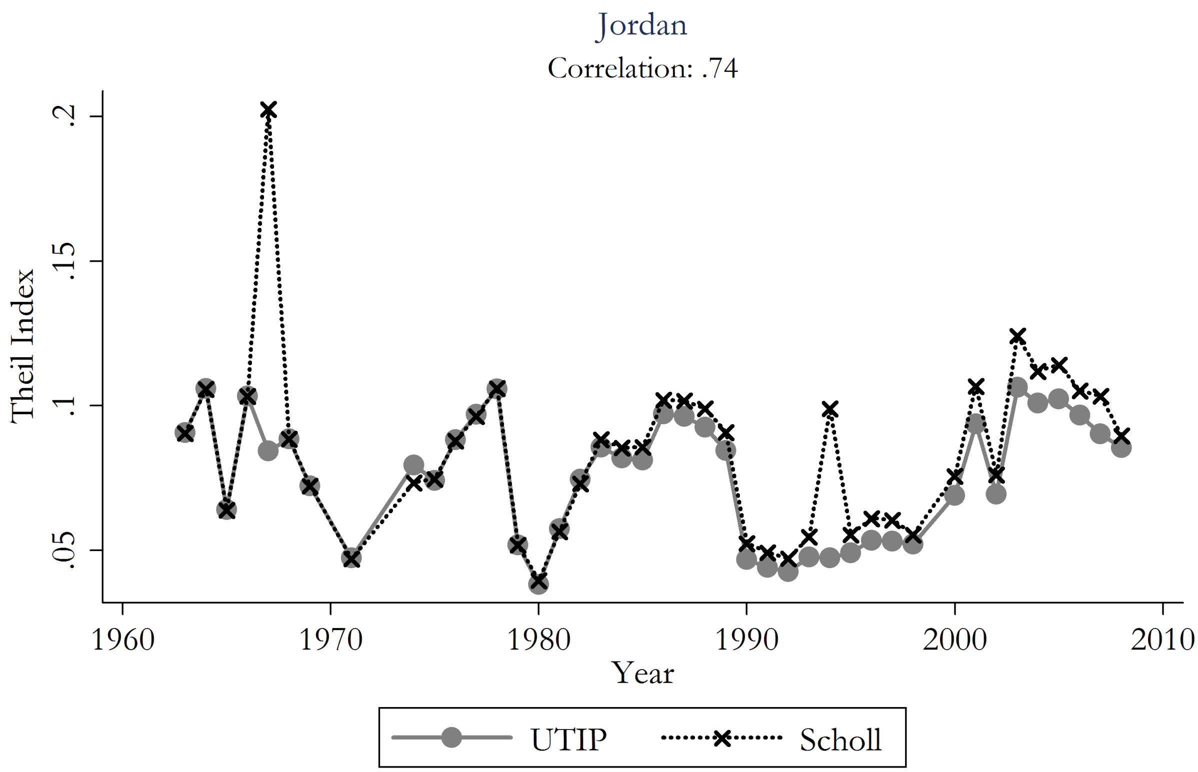
Jordan
Correlation: .74
Theil Index
.05
.1
.15
.2
1960
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A25

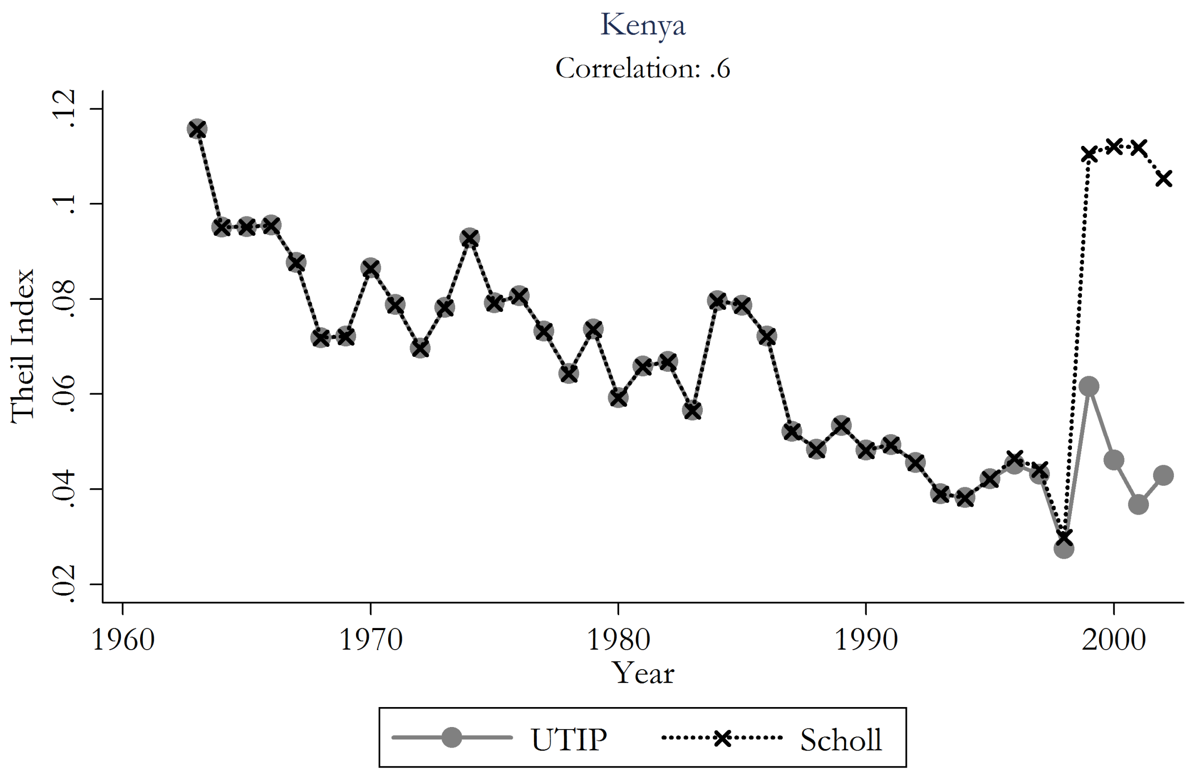

Kenya
Correlation: .6
Theil Index
.02
.04
.06
.08
.1
.12
1960
1970
1980
1990
2000
Year
UTIP
Scholl

Figure A26

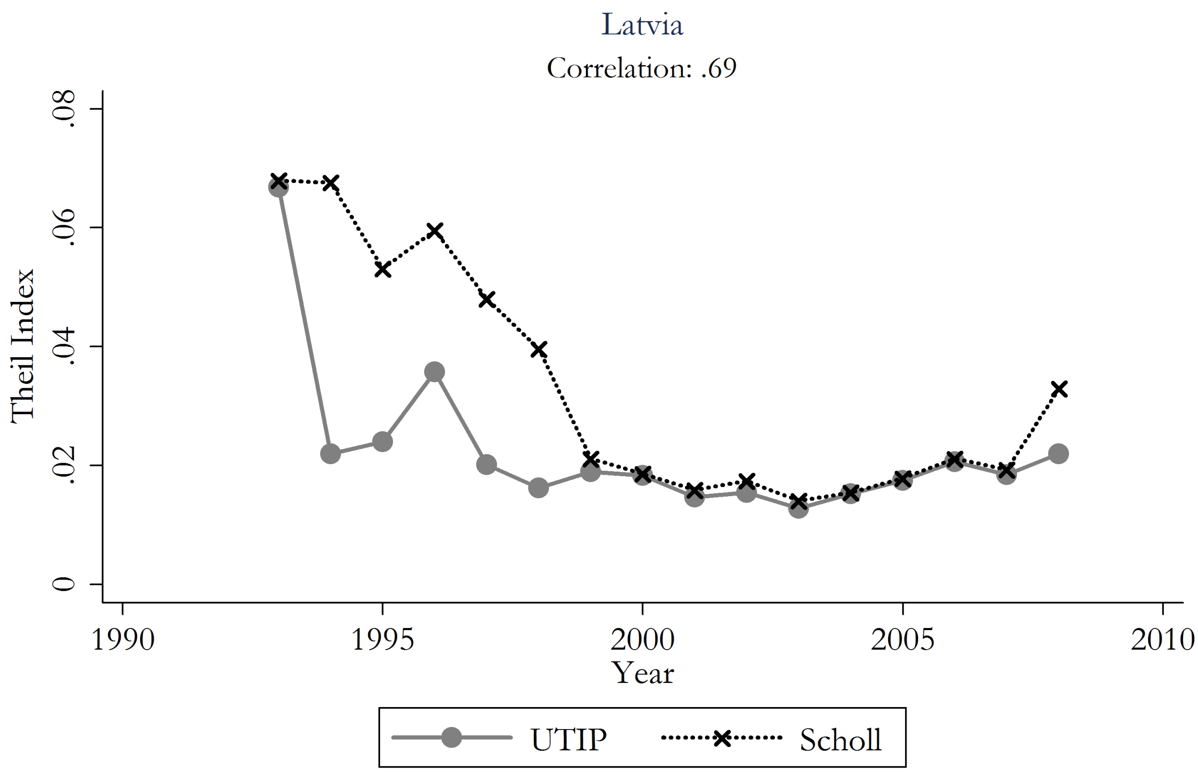
Latvia
Correlation: .69
Theil Index
0
.02
.04
.06
.08
1990
1995
2000
2005
2010
Year
UTIP
Scholl

Figure A27

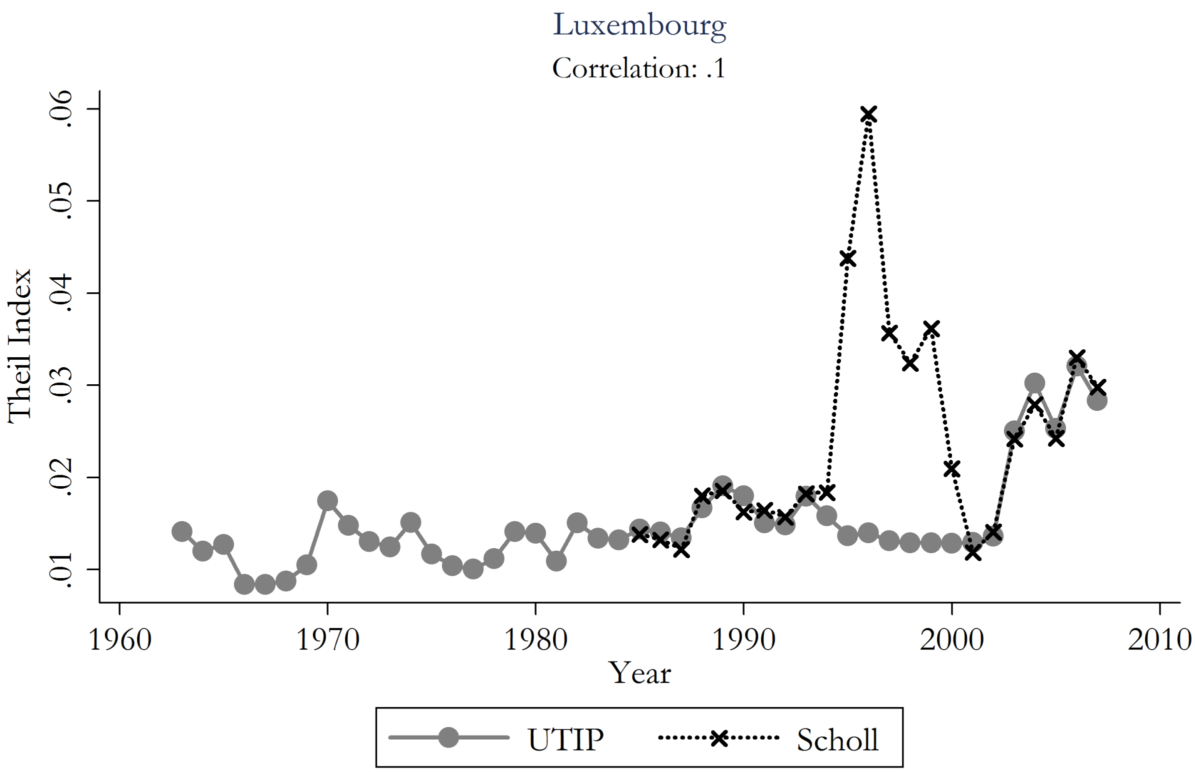

Luxembourg
Correlation: .1
Theil Index
.01
.02
.03
.04
.05
.06
1960
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A28

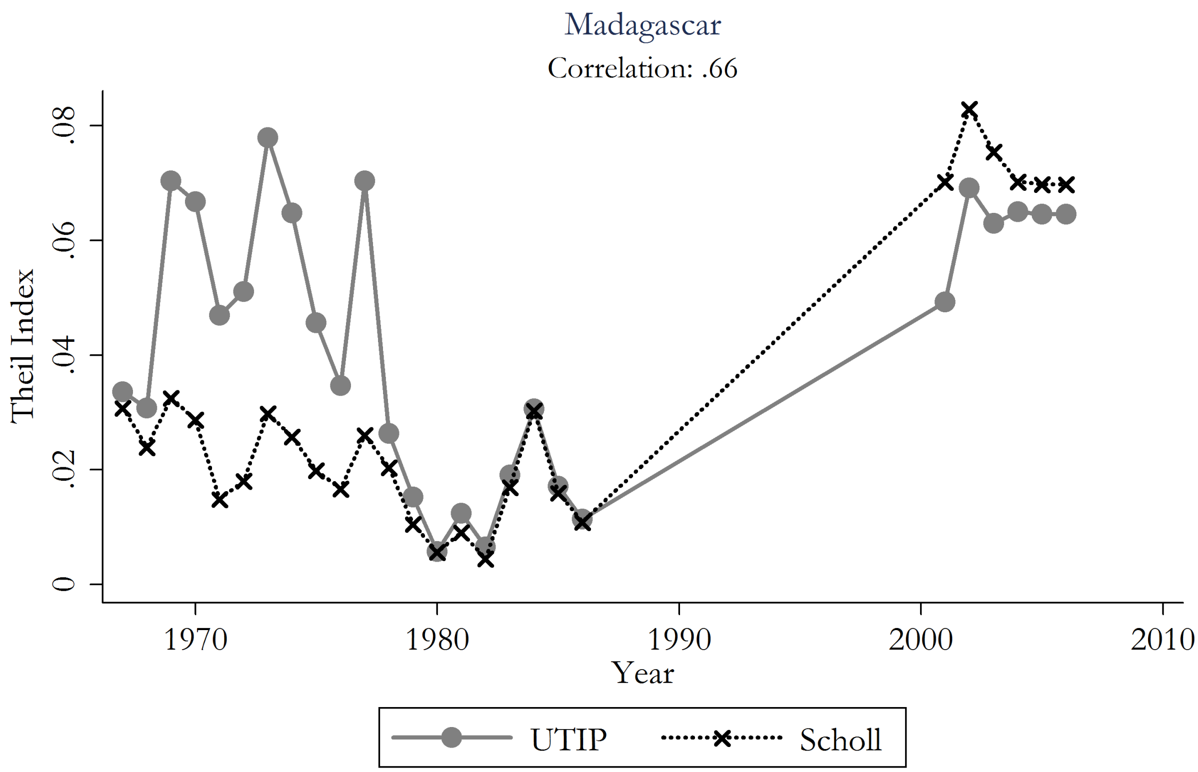

Madagascar
Correlation: .66
Theil Index
0
.02
.04
.06
.08
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A29

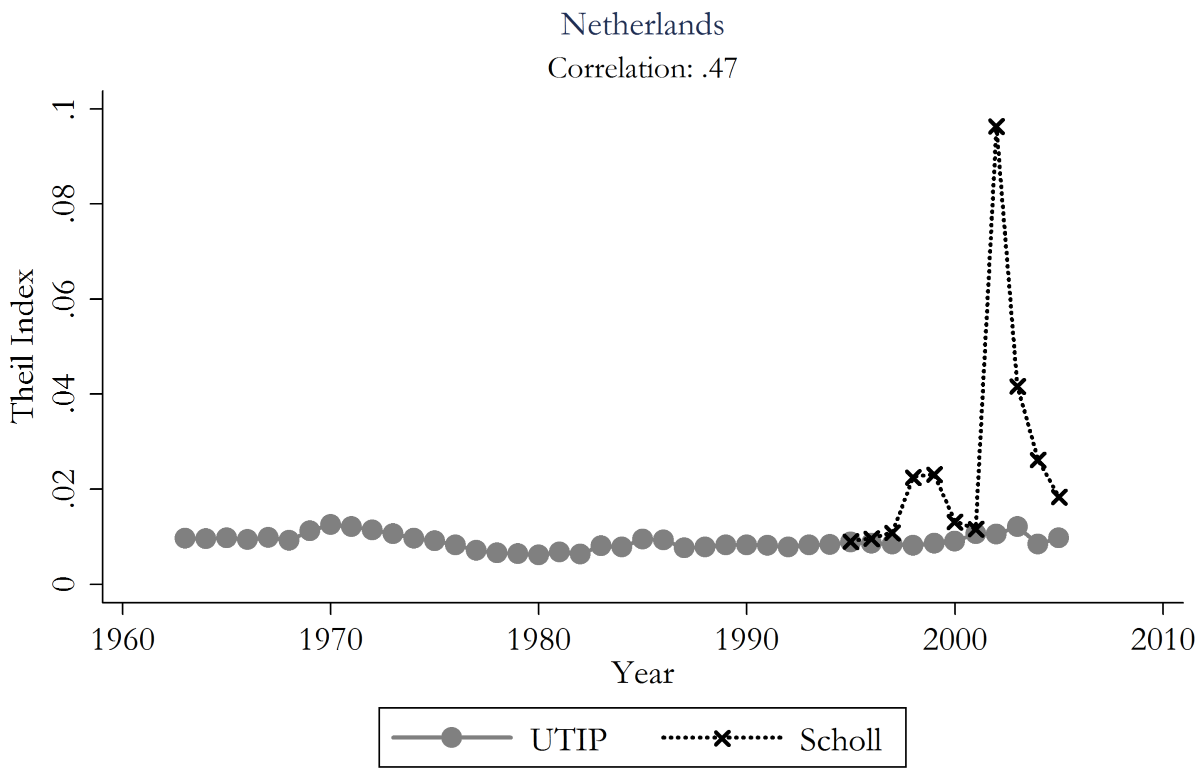
Netherlands
Correlation: .47
Theil Index
0
.02
.04
.06
.08
.1
1960
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A30

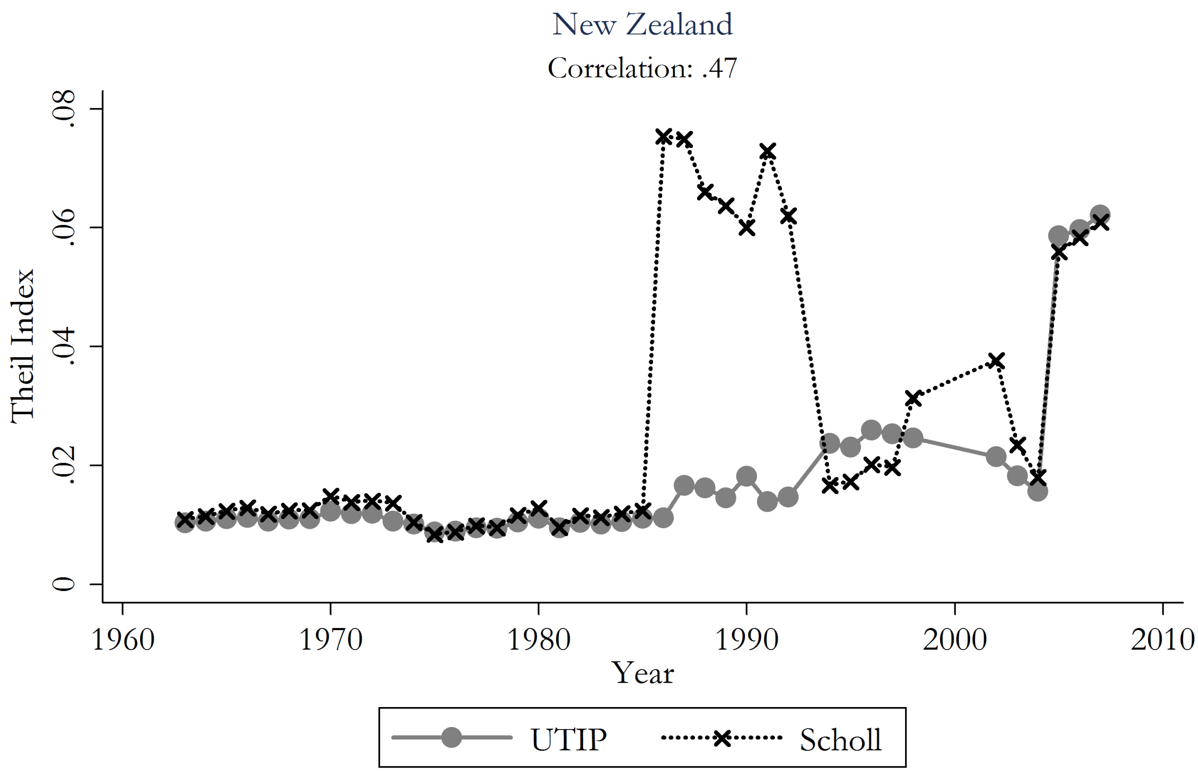

New Zealand
Correlation: .47
Theil Index
0
.02
.04
.06
.08
1960
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A31

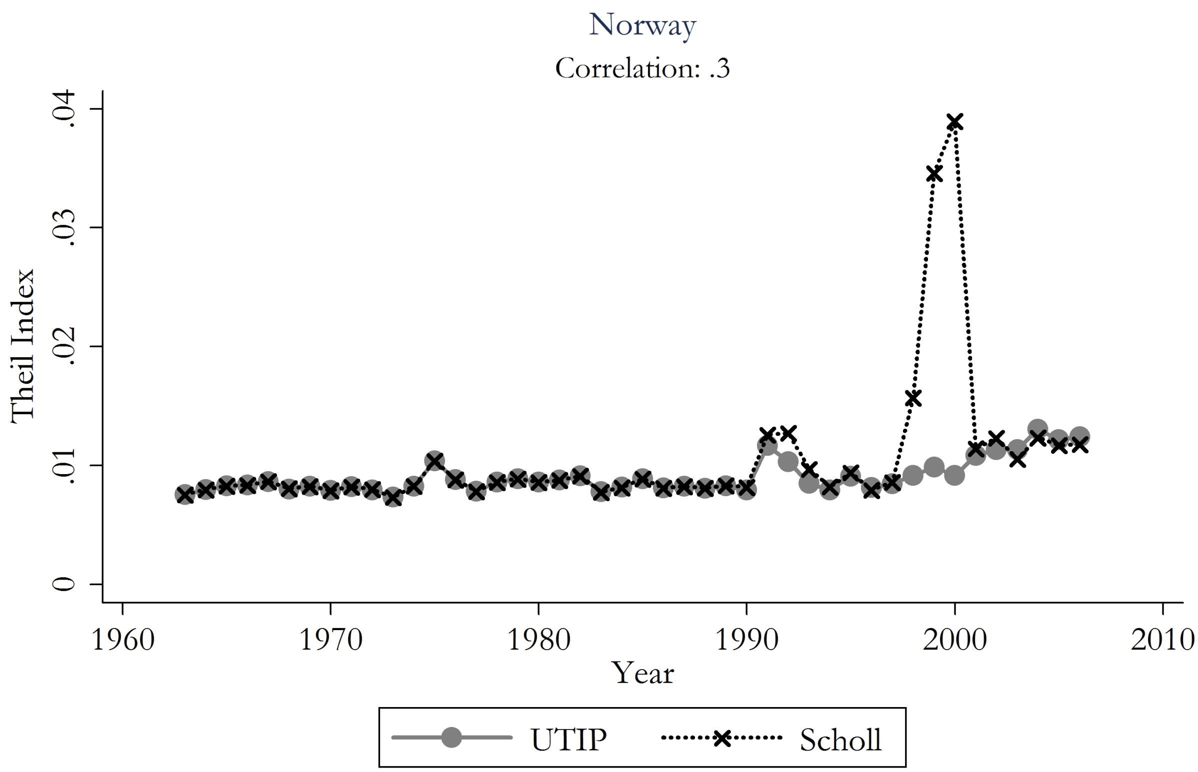
Norway
Correlation: .3
Theil Index
0
.01
.02
.03
.04
1960
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A32

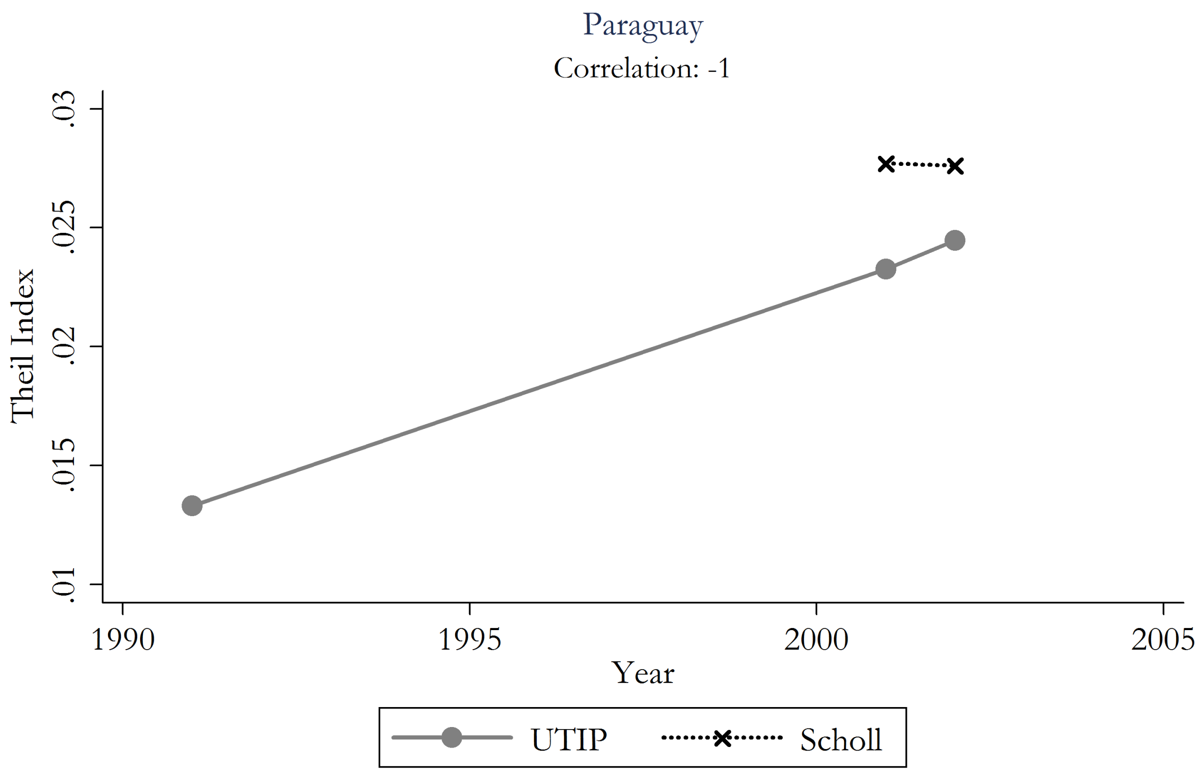
Paraguay
Correlation: -1
Theil Index
.01
.015
.02
.025
.03
1990
1995
2000
2005
Year
UTIP
Scholl

Figure A33

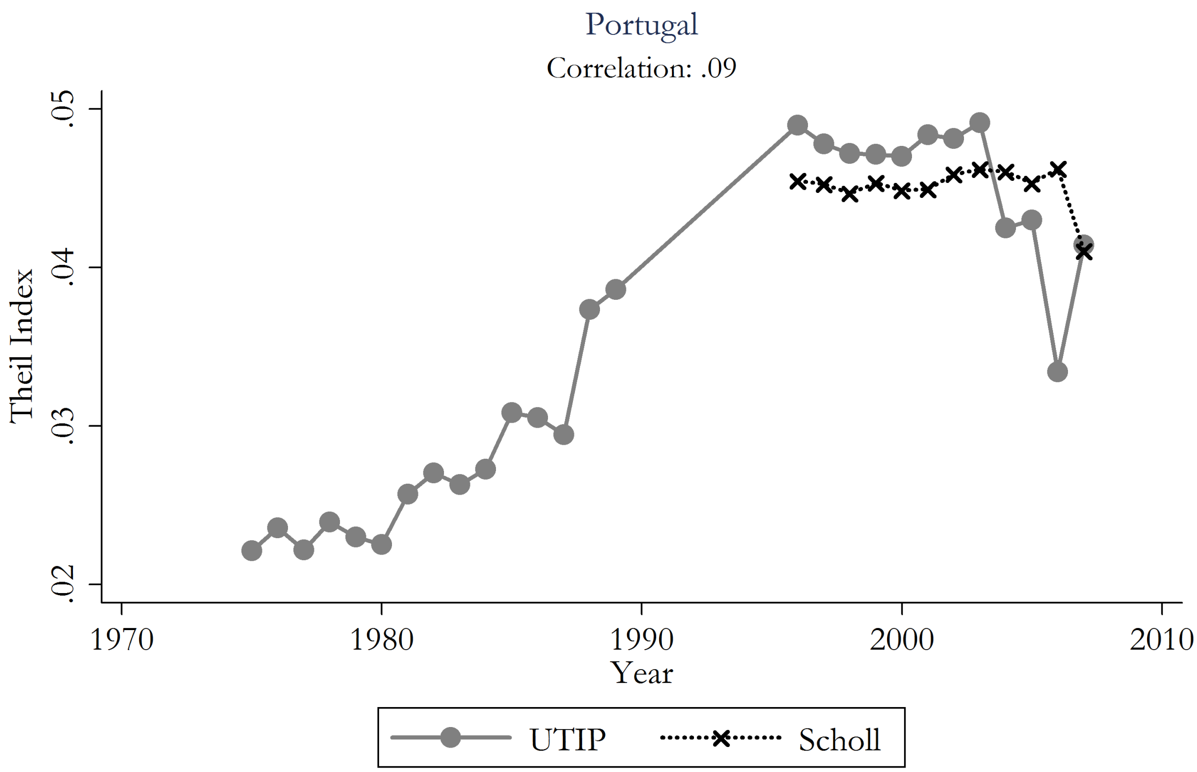
Portugal
Correlation: .09
Theil Index
.02
.03
.04
.05
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A34

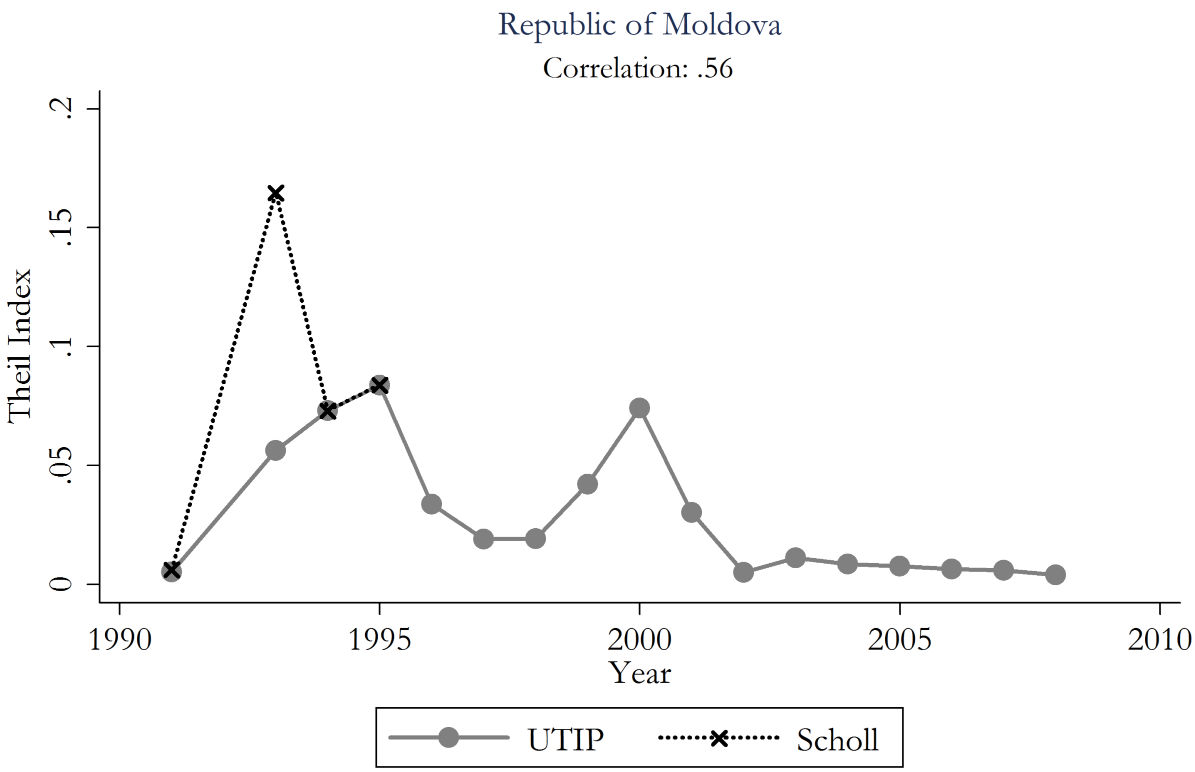
Republic of Moldova
Correlation: .56
Theil Index
0
.05
.1
.15
.2
1990
1995
2000
2005
2010
Year
UTIP
Scholl

Figure A35

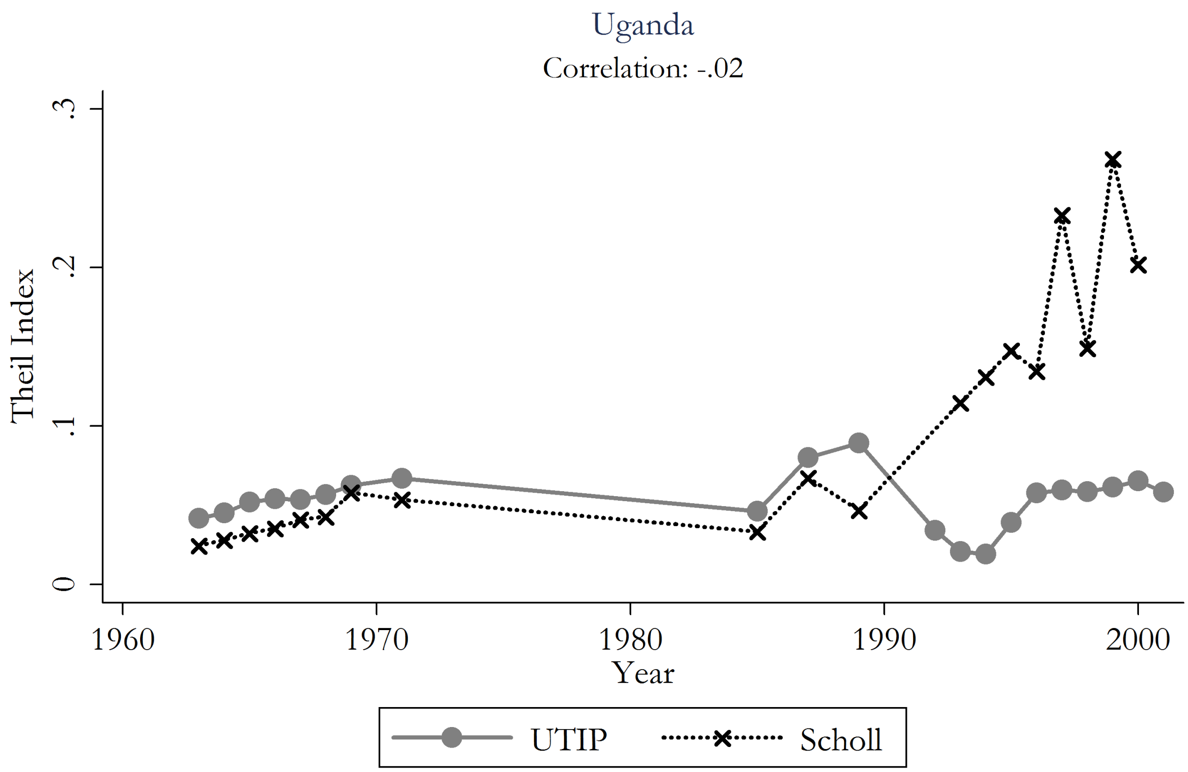

Uganda
Correlation: -.02
Theil Index
0
.1
.2
.3
1960
1970
1980
1990
2000
Year
UTIP
Scholl

Figure A36

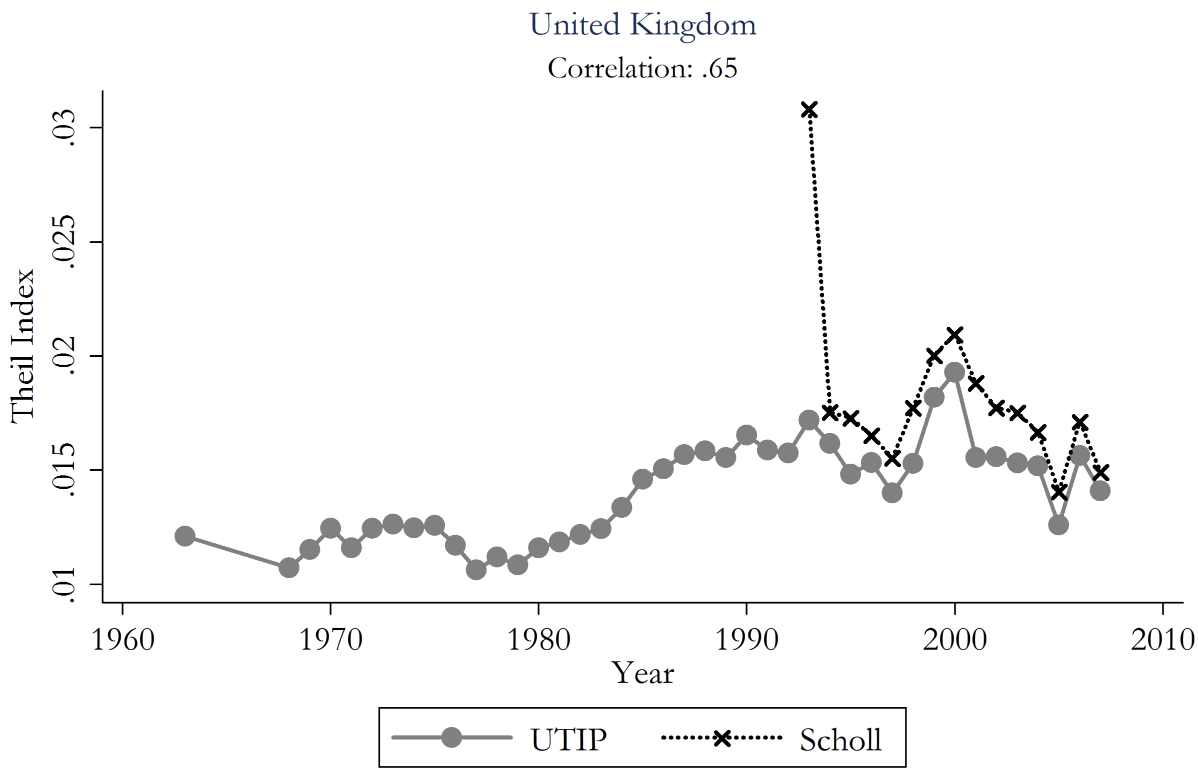

United Kingdom
Correlation: .65
Theil Index
.01
.015
.02
.025
.03
1960
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A37

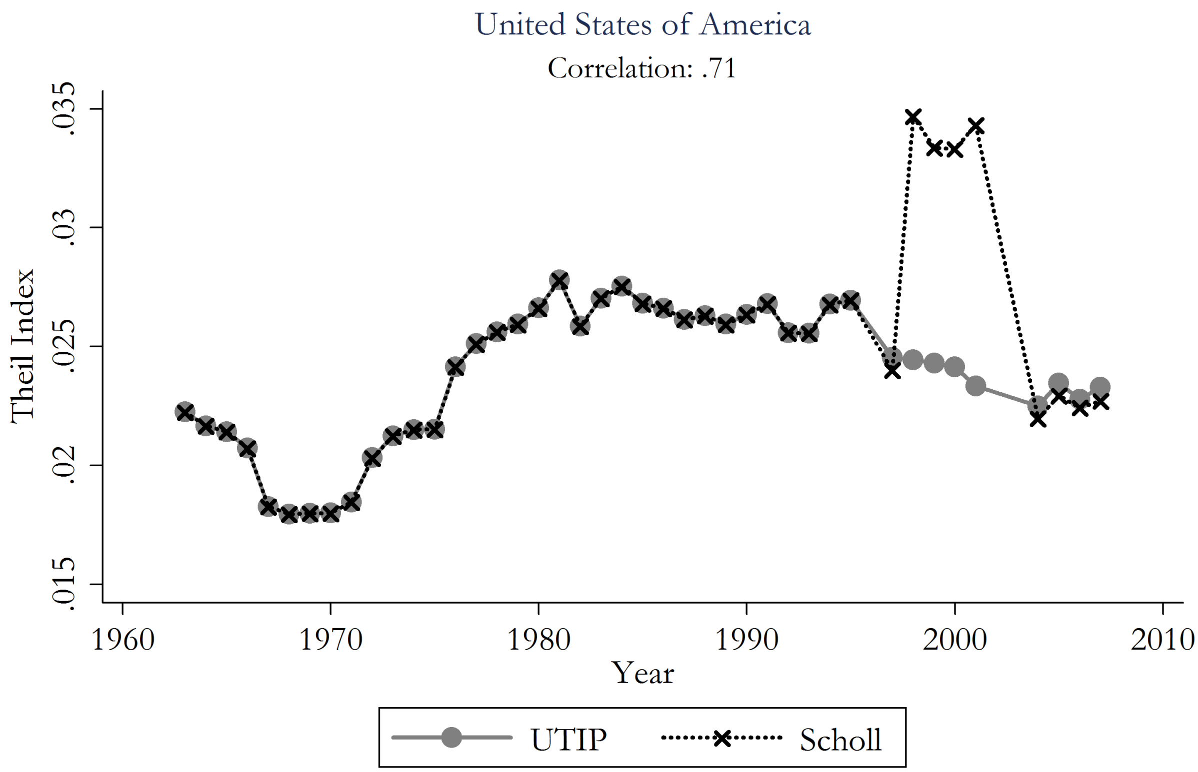
United States of America
Correlation: .71
Theil Index
.015
.02
.025
.03
.035
1960
1970
1980
1990
2000
2010
Year
UTIP
Scholl

Figure A38

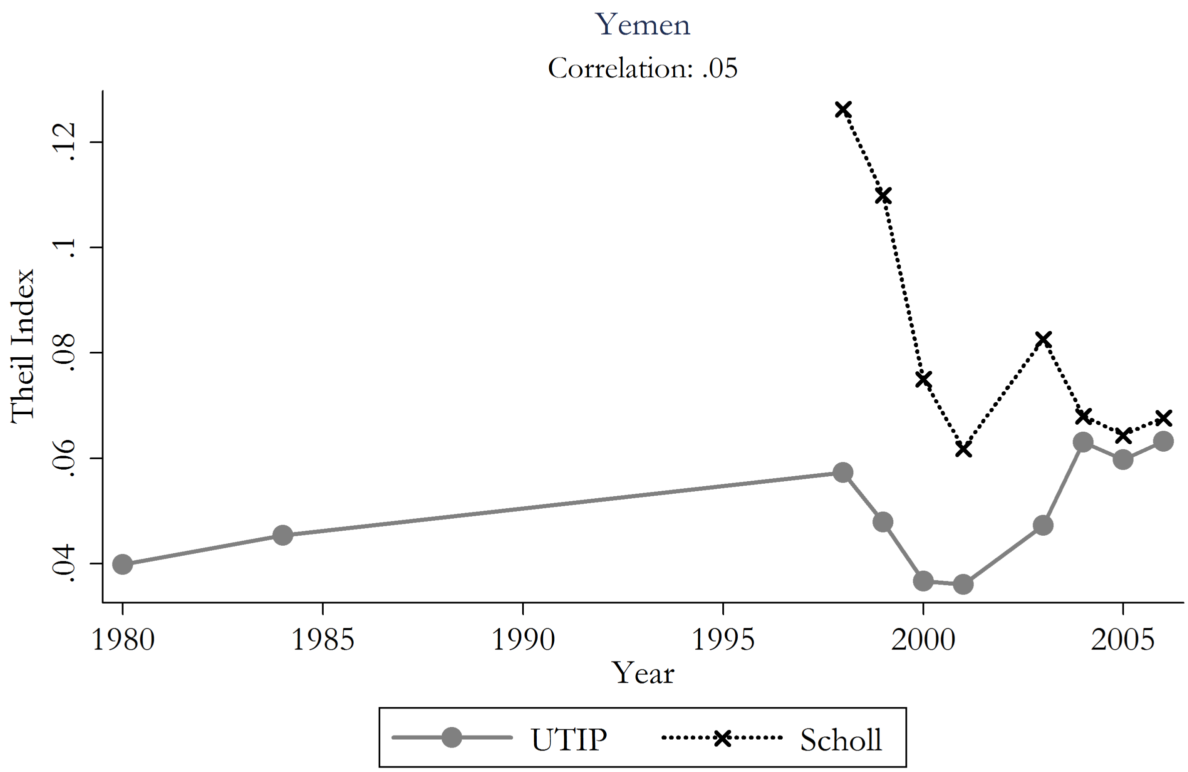

Yemen
Correlation: .05
Theil Index
.04
.06
.08
.1
.12
1980
1985
1990
1995
2000
2005
Year
UTIP
Scholl