



LUND UNIVERSITY

Estimating Historical Inequality from Social Tables: Towards Methodological Consistency

von Fintel, Dieter; Links, Calumet; Green, Erik

2023

Document Version:
Other version

[Link to publication](#)

Citation for published version (APA):

von Fintel, D., Links, C., & Green, E. (2023). *Estimating Historical Inequality from Social Tables: Towards Methodological Consistency*. (pp. 1-36). (Lund Papers in Economic History.; No. 2023:247).

Total number of authors:
3

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Lund Papers in Economic History



No. 247, 2023

Estimating Historical Inequality from Social Tables: Towards Methodological Consistency

Dieter von Fintel, Calumet Links, Erik Green

Estimating Historical Inequality from Social Tables: Towards Methodological Consistency*

Dieter von Fintel^{†1,2}, Calumet Links^{‡1,2}, and Erik Green^{§3,2}

¹*Department of Economics, Stellenbosch University, South Africa*

²*Laboratory for the Economics of Africa's Past, Stellenbosch University, South Africa*

³*Department of Economic History, Lund University, Sweden*

Abstract

Research on long-term historical inequality has expanded to include previously neglected periods and societies, particularly in the global South. This is partly due to the resurgence of the social tables method in economic history, an approach which uses archival records to reconstruct income and wealth distributions in contexts where micro data is unavailable. This method can cause a downward bias in estimating inequality, but there is limited evidence of this bias in economic history. We collected a new data set of 108 historical social tables spanning over a 1000 years. We found that the compilers consistently made careful methodological choices that took data limitations into account. We found that the inequality estimates are not systematically related to the number of classes chosen or the size of the top class, but that choosing bottom classes that bundle together even small variations in income or wealth can introduce a downward bias to the inequality estimates. This drawback can be overcome by using methodological cohesion to mitigate the problem of limited information about the poorest classes in colonial archives.

Keywords: Social tables, Gini, inequality, pre-industrial, grouped data

*This paper draws data from the work of many other researchers, including the significant contributions of Branko Milanovic and his colleagues, as well as the African Long-term Inequality Trends (AFLIT) research group. This analysis would not have been possible were it not for the archival and transcription efforts of these groups of researchers. All errors, of course, remain those of the authors.

[†]dieter2@sun.ac.za

[‡]calumetlinks@sun.ac.za

[§]erik.green@ekh.lu.se

I Introduction and overview

I.I Understanding long-run inequality

Pre-industrial economic inequality has attracted attention in the recent economic history literature, as new data sources enable researchers not only to document early inequality but also to investigate the origins of inequality today (Alfani 2021; Milanovic 2018; Milanovic 2009; Milanovic 2006; Berry 1993; Booth 1988; Fourie and Fintel 2010; Fourie and Fintel 2011). Economic historians have focused particularly on the growing inequality during and after industrialization. Kuznets (1955)’s seminal work hypothesizes an inverted U-shaped relationship between long-term economic growth and inequality, with low inequality in the pre-industrial era, an increase in inequality during industrialization, and eventually a reduction after a certain threshold is reached. Recently, Piketty (2018), Atkinson, Piketty, and Saez (2011), Milanovic (2006), and Milanovic (2022) have also contributed to our understanding of how inequality evolves over the long run and how it relates to economic development. Kuznets (1955), on the one hand, claims that inequality was relatively low before the industrial revolution; Van Zanden and Van Leeuwen (2012) and Van Bavel and Van Zanden (2004) on the other, show that income inequality was quite high in some areas of Europe over the same period. Despite this surprising finding, Van Zanden and Van Leeuwen (2012) nevertheless consider the possibility of a “super-Kuznets curve” emerging in Holland. They hypothesize that the origins of an inverted U-shaped Kuznets curve may lie in the pre-modern growth period (before 1800) and not in the advent of the industrial revolution (Van Zanden and Van Leeuwen 2012).

In a recent meta-study, Alfani (2021) combined pre-industrial inequality studies mainly from Europe with some from non-European areas (such as Ottoman Anatolia, pre-revolutionary USA, and Tokugawa Japan)¹. He concludes that the dramatic increase in the number of inequality studies since 2010 has redefined the way economic historians look at long-term trends in inequality. The accepted view that long-run economic growth and inequality demonstrate an inverted U relationship has been extensively disputed, forcing researchers to look for alternative mechanisms to explain why pre-industrial inequality was high (Alfani 2021, p. 5). Population growth, for example, may have led to increased inequality. This is consistent with a Malthusian hypothesis by which finite natural resources are initially shared relatively equally between few people. In time, if populations grow faster than technological innovation, finite resources become more concentrated in the hands of elites. High population growth and inequality would place stress on natural resources leading to starvation, famine and disease, reducing population numbers once more. Another alternative explanation for high pre-industrial inequality in some European states may be the regressive tax systems that were implemented by political elites to favour their own interests and maintain their status, wealth and power. (Alfani 2021; Milanovic 2018).

Methods of estimation have played a crucial role in expanding the scope of research on historical inequality. Many studies have utilized social tables or data grouping techniques to calculate inequality metrics for both pre- and post-industrial societies, as individual or household level data is not always readily available (Van Ourti and Clarke 2011; Shorrocks and Wan 2008; Milanovic, Lindert, and Williamson 2007). Researchers use groups of people that are broadly representative of the socio-economic spectrum of a society, even if the individuals in those groups are not as well-documented as in census data or modern household surveys. Archival sources, such as government records and other accounts, are used to assign aggregate income or wealth values to the groups and to determine group sizes for generating a population distribution. The

¹Many of the early modern era inequality studies referenced in Alfani (2021), such as those for Japan and the United States, do not use social tables but rely on micro records from population censuses compiled for administrative purposes.

socio-economic differences between these groups are assumed to provide a reliable representation of overall inequality, despite some differences within the groups remaining unaccounted for. Recent studies on long-term inequality have been made possible by transcribed archival data and the creation of social tables from these sources (Milanovic 2006; Milanovic 2018; Milanovic 2009; Milanovic, Lindert, and Williamson 2011). This has led to new understanding of the connections between long-term economic development, institutions, globalization, and inequality (Bolt and Hillbom 2016; Saito 2015; Milanovic, Lindert, and Williamson 2011).

In this study, we investigate how the social tables methodology has contributed to research on long-term global inequality. We identify specific biases that may compromise the use of social tables in comparative economic history, focusing on the limiting effect of using large subsistence classes in constructing Gini coefficients. Additionally, we systematically analyze how various estimation factors, particularly the size of the bottom class, affect social table inequality estimates. We establish the conditions under which a large bottom class leads to the underestimation of inequality. We find that other potential weaknesses of social tables are less serious. We conduct a meta-study of 108 social tables, complementing the work of Milanovic, Lindert, and Williamson (2011) and Milanovic (2018) by including newer studies published between 2018 and 2021. Our results show that the size of the bottom class is significantly negatively related to inequality estimates in populations smaller than two million.

I.II Opportunities in using social tables

Measuring inequality in comparable ways across countries and periods is crucial for the growing literature on long-run global inequality. However, it is difficult to maintain stable measurement concepts and data sources to make credible comparisons. This paper looks at how the social tables method has contributed to this literature.

Social tables are used to measure inequality when micro records are not available and researchers are obliged to reconstruct historical income and wealth distributions using limited group-level data (Van Ourti and Clarke 2011; Cowell 1991; Shorrocks and Wan 2008). The first social table was constructed by Gregory King in the 1930’s to map out the incomes and social structure of England and Wales in 1688. The approach, subsequently popularized by Milanovic, Lindert, and Williamson (2011), has contributed to new understanding of the connections between long-term economic development, institutions, globalization, and inequality (Bolt and Hillbom 2016; Saito 2015; Milanovic, Lindert, and Williamson 2011; Sokoloff and Engerman 2000).

Many studies have used techniques like social tables or data grouping in order to calculate inequality metrics for pre-industrial societies because individual or household level data is not readily available. Researchers use groups that are broadly representative of the socio-economic spectrum of a society, even if the individuals in those groups are not as well documented as those in census data or modern household surveys. Archival sources, such as government records, are used to assign aggregate income or wealth values to the groups and to determine group sizes for generating a population “distribution”. The socio-economic differences between these groups are assumed to provide an accurate representation of overall inequality, despite some differences within the groups remaining unaccounted for.

Grouping data is useful not only in historical studies but also in present-day research where, to protect people’s privacy, individual data is not disclosed (Van Ourti and Clarke 2011; Cowell 1991; Shorrocks and Wan 2008). Various statistical techniques have been developed to handle grouped data and these are widely discussed in the literature. However, apart from Modalsli (2015), economic historians have not extensively examined the limitations of grouped data.

The use of social tables based on archival sources to estimate inequality in various societies and historical

periods, originally popularized by Milanovic, Lindert, and Williamson (2011), has reawakened interest in the origins of inequality both historically and today. This method allows us to study regions and societies where data is sparse, such as those in the global South, filling knowledge gaps and altering our perspectives on how inequality relates to economic development². A recent survey by Galli, Theodoridis, and Rönnbäck (2022) of the literature on long-term inequality trends in Africa and Latin America shows that such studies are expanding into regions of the world that have traditionally received very little attention.

Along with the increased interest in long-term inequality, over the past decade the number of studies using social tables to measure inequality has multiplied. Figure 1 shows that social tables took off shortly after the appearance of the first working paper by Milanovic, Lindert, and Williamson (2007), which popularized the method. Up to that point, social tables had been constructed mainly for Asia and Europe, with fewer for Africa and the Americas; afterwards, the emphasis was on South America and Africa, where inequality levels are high today. Figures 2 and 3 show that social tables used by Milanovic, Lindert, and Williamson (2011) and Milanovic (2018) dominated the studies of the pre-industrial period (pre-1800)³. While their initial collection also included some studies of the “post-industrial” period (post-1800), it was only after 2008 that other researchers added a significant number of these to the literature⁴. These patterns show that there was little research – for either pre- or post- industrial societies – using social tables before the work of Milanovic, Lindert, and Williamson (2007), but that the method they used for pre-industrial societies was rapidly adopted for studies of post-industrial societies in which micro data were not available for estimating inequality, thus contributing more than simply an understanding of “ancient inequality”.

Despite the progress made in expanding the evidence base using social tables, we should not overlook the importance of testing the validity of the findings from these new studies. In the rest of the paper we investigate possible sources of bias in social tables and evaluate their impact through a new meta-study.

I.III Cautions about using social tables

Although grouping methods are currently the best way to analyse inequality in contexts that lack micro data, their limitations must be acknowledged at the start. Researchers need to be aware of them so they can determine whether their inequality estimates can be accurately compared across time and countries.

I.III.1 Within-group inequality

Milanovic, Lindert, and Williamson (2007) anticipated one of the major limitations. Ignoring within-group variation is equivalent to calculating the Gastwirth (1972) “lowerbound” estimate of inequality from grouped data. Most historical social table studies rely implicitly on this method, and therefore tend to underestimate inequality. To circumvent this, Milanovic’s (2018) comparative study uses micro data inequality estimates as far as possible – including those from Holland in 1526, Tuscany in 1427, Old Castille in 1752 and the Netherlands in 1808. Other studies that use social tables do not address this concern. Micro data from the Kingdom of Naples in 1811 and Bihar India in 1807 reveal that within-group inequality Gini estimates

²For example, the African Long-Term Inequality Trends (<https://www.afil.net/>) project is reconstructing a wide range of inequality estimates for data-poor countries in Africa that have high levels of inequality today. Using social tables could help explain how these societies became like this.

³We roughly define the pre-industrial period in Europe and most of the current global North as pre-1800, and the period when these countries had become industrialized as post-1800. Some countries in the current global South were of course still only beginning to industrialize in the post-1800 period. We use “post-industrial” as shorthand for “post-industrialization”

⁴A large number of social tables were generated in 2008. These come from one study on Latin American inequality by (FitzGerald 2008). We show these numbers separately. . We show these numbers separately.

were high – above or close to 0.5 (Williamson 2010; Lindert and Williamson 1982; Bértola et al. 2010). The problem is exacerbated in years where data is scant or of poor quality and may engineer what Milanovic (2018) calls “Kuznets-wave-like periods of waxing and waning inequality”.

Milanovic, Lindert, and Williamson (2007) also mitigate these concerns by estimating the “upperbound” estimate proposed by Gastwirth (1972) for all the available social tables. The adjustment allows us to account for the maximum possible variation within groups. Average differences between conventional and adjusted Gini coefficients are, according to Milanovic, Lindert, and Williamson (2007), not serious enough for further concern.

I.III.2 Overlapping Classes

Modalsli (2015) adds a further caution as regards social tables. Ignoring the fact that classes overlap in reality may lead researchers to underestimate inequality. For example, a shopkeeper or artisan in 1788 France might have enjoyed the same income as a noble or a cleric. The social table constructed by Morrisson and Snyder (2000) assumes that, on average, all nobles and clerics were five times richer than all shopkeepers and artisans. This issue of overlapping classes may be the most difficult to get round.

As Kakwani’s (1980) geometric exposition shows, the Gastwirth (1972) correction relies on using non-overlapping, mutually exclusive classes to estimate inequality from grouped data. But Modalsli (2015) shows that even the Gastwirth correction does not go far enough. The “simple” correction proposed by Van Ourti and Clarke (2011) also does not correct for overlapping classes. Modalsli (2015) proposes using parametric distributional assumptions that simultaneously model within-group inequality and allow groups to overlap. But he does not give guidance on the extent of class overlap, and it is therefore hard to establish which assumptions are appropriate. His approach does not offer unique upper bounds Gastwirth (1972), but it does emphasize that there are circumstances in which inequality estimates are understated because of overlapping classes. The extent of this bias remains unquantifiable. This downward bias, whose extent remains unquantifiable, remains a problem when comparing trends in inequality over the long run and across regions – especially when the bias is context-specific.

I.IV Number of classes

One way to minimize both within-group and overlap bias is to use a larger number of classes. With a small number of classes, the population size of each class is large and there is little between-group variation in imputed income or wealth. As we increase the number of classes, the population size of each class becomes smaller, between-group variation increases, and the data evolves naturally into micro-data that producing a finer distribution of the entire population. Examples of social tables with very few classes are Moghul India in 1750, with a population of 182 million and only four classes, and China in 1880, with a population of 377 million people and only three classes (Milanovic, Lindert, and Williamson 2011). Both Milanovic (2018) and Modalsli (2015) conclude that changing the number of classes does not significantly predict inequality estimated from social tables. Their findings imply that researchers should not worry unduly about this possible problem. However, both studies use relatively few inequality estimates to reach their conclusions. We argue that their estimates have insufficient statistical power to reveal a relationship between the number of groups and the accuracy of the inequality estimates, and further that their limited sample omits context-specific factors. We address both concerns by expanding the number of studies in the social tables database and re-assessing this limitation.

I.IV.1 The size of the top class

Even if having few classes is ruled out as a source of bias, it may still be a problem because it invites another source of bias in the relative size of the classes. Two concerns are an over-large top class and an over-large bottom class. Either can happen when too few classes are chosen for the analysis. These two kinds of disproportion are specific manifestations of what has traditionally been considered a more general problem.

We start with the first. Alvaredo (2011) shows that the Gini coefficient directly relates to the share of income or wealth held by an “infinitesimal” share of the population at the top of the distribution. The whole-population Gini coefficient is a function of the top class’s share in income or wealth (S) and the Gini coefficient within the rest of the distribution (G^*):

$$Gini = G^*(1 - S) + S \tag{1}$$

Alvaredo’s formulation shows that if S is highly concentrated, benefiting only a small proportion of the population, and a large majority of the population is characterized by low inequality (G^*), the share of income or wealth held by the top class can dominate total Gini inequality.

This relationship has two implications. Focusing only on the “very top” income or wealth shares can reveal trends in overall inequality. Indeed, the recent literature on inequality has put a great deal of effort into monitoring top income and wealth shares (Alvaredo et al. 2013; Atkinson 2007; Kopczuk 2015; Saez and Zucman 2016; Piketty, Saez, and Zucman 2022). But the second implication is more relevant here. The very top part of income and wealth distributions usually follows an extreme value distribution. Typically, survey data is inadequate for measuring the full extent of the skewness in top tails, while tax records fare better. Estimating the full extent of inequality therefore means we have to impute “high” incomes for “few” people in the extreme top tails, using Pareto distributions (Pareto 1897; Jenkins 2017; Hundenborn, Woolard, and Jellema 2019; von Fintel and Orthofer 2020). Accounting for these extremes has a bearing on inequality estimates, and shows how sensitive overall inequality is to inequality at the top.

The very distinct nature of the extreme top tail implies that the top classes in social tables should be as small as possible – or even “infinitesimal” – and should represent only the elite of the society. However, social tables are often unable to distinguish between the “very rich” and the “super rich”, and they bunch these arguably distinct groups together. Grouping them together could introduce downward bias to inequality estimates. Modalsli’s (2015) finding that – as with the number of classes – there is no relationship between the proportion of the population in the top class and social table inequality estimates therefore comes as a surprise. Again, his null result may be the result of having insufficient statistical power to detect a significant relationship. While one could argue that historical societies do not require the same distinctions within the top tail of the distribution, Pareto’s early work (1897) suggests otherwise (Pareto 1897).

I.IV.2 The size of the bottom class

An additional concern raised by Modalsli (2015), but which he does not investigate directly, is the impact of the size of the bottom income group or class on inequality estimates. In contrast to the top tail, the bottom tail of income and wealth distributions typically contain high concentrations of *similarly* poor individuals, with low variation within this group. Grouping at the bottom should, presumably, have a smaller impact on bias in inequality estimates than grouping at the top. It is therefore not immediately apparent that social tables researchers should prioritise dividing the bottom tail into smaller classes over introducing granularity in the top class. This presumption is the likely reason why researchers typically do not raise this concern. However, we argue that constructing social tables nonetheless requires using relatively small *bottom* classes.

An additional concern raised by Modalsli (2015), but which he does not investigate directly, is the way the size of the bottom income group or class affects inequality estimates. In contrast to the top tail, the bottom tail of income and wealth distributions typically contains large numbers of similarly poor individuals, with little variation between them. This being the case, the bottom group should presumably be less likely to cause bias in inequality estimates than the top group. It is therefore not immediately apparent that social tables researchers should be more concerned to have smaller classes in the bottom group than the top. We argue, however, that social tables nonetheless require relatively small bottom classes.

Obtaining detailed information at the bottom of the income distribution from archival sources has typically been difficult, especially since the poor tend to work outside the formal economy (Aghevli and Mehran 1981). In Africa, for example, there are no official historical statistics on subsistence production and the degree of by-employment. This is because most of the economic activities of the poor – past and present – occur outside the ambit of the formal economy. Aware of this fact, some researchers who have used social tables (Aboagye and Bolt 2021; Bolt and Hillbom 2016; Haas 2017) have been careful to impute reliable, context-specific mean subsistence incomes to bottom classes. However, researchers have paid little attention to reducing the size of the bottom class – a failing often attributed to not having sufficient archival information. A recent study on inequality in colonial Ghana (Aboagye and Bolt 2021), classifies between 37 and 60% of the population as subsistence farmers. In 1861 Chile, 40% of the population are classified as “fishermen” or the bottom class (Milanovic 2018).

Those examples illustrate the difficulty of identifying small differences in income or wealth in the bottom group. When this group constitutes a large proportion of the population, inequality estimates may be sensitive to these small differences.

We consider the problem of a large bottom class more systematically in appendix A. Our starting point is Dagum’s decomposition of the Gini coefficient into within-group and between-group components, and we follow reasoning similar to Alvaredo’s (2011) (Dagum 1997). Our derivations express inequality estimates in terms of observable social table attributes (class proportions and means). We first analyse the hypothetical scenario “ L ”. It uses a social table with k classes, a “large” bottom class with population share P_1^L and unobserved “within-class” Gini of G_{11}^L . The same social table is re-analysed in scenario “ S ”, but the bottom class is sub-divided into two smaller groups. This yields $k + 1$ classes, with the bottom group proportions adding up to the “large” bottom class ($P_1^S + P_2^S = P_1^L$). The two sub-divided groups in the bottom tail have unobserved within-group Gini coefficients G_{11}^S and G_{22}^S . We derive a range of stylised facts by analysing the difference in social table Gini estimates under the two scenarios ($\widehat{G}_S - \widehat{G}_L$).

Firstly, equation 2 shows that $\widehat{G}_S > \widehat{G}_L$. This confirms that inequality is underestimated with larger bottom classes. The magnitude of the change is determined by the differences in the means between the two new, smaller bottom classes. If μ_2^S is substantially larger than μ_1^S , adding another class increases the Gini estimate considerably. That is, separating sub-populations that are materially different from each other has an important consequence for estimating inequality from social tables. This is of course easier said than done: first we have to overcome the difficulty of identifying groups with distinct socioeconomic status in the bottom class. However, it is worth the effort, as even small mean differences can have a meaningful impact on the Gini estimate. If either of the sub-divided groups has a large proportion in the total population ($P_1^S; P_2^S$), small differences in their average incomes can make a big change in inequality estimates. This is a more realistic scenario for social tables researchers. But even after gathering exhaustive archival evidence to identify differences, the bottom class may still be large, and have big enough within-group differences to introduce bias into estimation.

$$\widehat{G}_S - \widehat{G}_L = P_1^S P_2^S \frac{\mu_2^S - \mu_1^S}{\mu^S} > 0 \text{ because } \mu_2^S > \mu_1^S \text{ and } \mu_1^S > 0; \mu_2^S > 0; P_1^S > 0; P_2^S > 0 \quad (2)$$

Secondly, we analyse changes in unobserved within-group inequality. Equation 3 verifies that $\widehat{G}_S > \widehat{G}_L$ after adding a class to the bottom of the social table, but provides additional insights into the reason for this change.

$$\widehat{G}_S - \widehat{G}_L = (G_{11}^L - G_{11}^S) P_1^S S_1^S + (G_{11}^L - G_{22}^S) P_2^S S_2^S + G_{11}^L (P_1^S S_2^S + P_2^S S_1^S) \quad (3)$$

Because the larger bottom class is divided into smaller classes, variation within the new groups is reduced, while variation across groups grows. It follows that $G_{11}^L - G_{ii}^S > 0 \dots j = 1; 2$, and this raises the overall Gini estimate. Because the first two terms are scaled by $P_j^S S_j^S \dots j = 1; 2$ (where S_j is the class share in total income), the change in within-group inequality is amplified by large population and income shares in the bottom classes. We are less concerned about large income shares, because the bottom of the distribution typically comprises groups that live in poverty. Rather, we pay more attention to large population shares in the bottom of the distribution, which is a typical limitation of many historical social tables. Where P_j^S relatively large, even small changes in within-group inequalities can have a large effect on the social tables inequality estimates. This finding is also reflected in equation 2. Large bottom class shares are particularly relevant in context where archives have only limited information for very large groups of informal workers or subsistence farmers, but where it may nevertheless be possible to trace many groups that have at least small variations in incomes.

We are even more concerned about large bottom classes when the total population size is small. Small populations typically have less diverse economies and are prone to being grouped into fewer, larger categories. However, as we argue, even small within-group differences can have a large effect on inequality estimates if the (bottom) class is too large. These factors are relevant for any study of pre-industrial inequality that uses social tables, but even more so for studies of countries in the global South that industrialized at a later date than the North (Milanovic 2006; Milanovic 2009; Milanovic, Lindert, and Williamson 2011; Milanovic 2018).

One way of assessing the robustness of social tables inequality estimates is by examining the stability of trends over time or using repeated studies of the same society and period. But this approach is complicated by the way a social table is constructed. For instance, the table may show the number and size of social classes changing across time, but this change may only be an artefact of the information contained in the archives and not a reflection of real economic change. For example, in the Ghanaian table the size of the bottom group declines from 60% in 1891 to 37% in 1960, with no identifiable economic cause associated with this change. This apparent reduction in the size of the bottom income group will inadvertently have an impact on the inequality estimates.

I.V The structure of the paper

In the preceding sections we noted that using social tables has been path-breaking for the field of economic history, but we also raised concerns about their statistical and theoretical underpinnings. We next provide empirical support for most of these concerns. We elaborate on the data we compiled for this meta-study and then use repeated case studies to make the case for possibly systematic relationships between inequality estimates and social table attributes. We next show how social table attributes evolved and infer that the

relationship between inequality estimates and these attributes is connected to a lack of data and not the result of any deliberate research design. We demonstrate that the size of the bottom class can bias the estimation of inequality in small populations. In conclusion we argue that methodological cohesion can mitigate the problem of limited information.

II Data

The data compiled for this study was collected from 2018 to 2022 and includes information from 108 social tables, a much larger number than the 41 initially analysed by Milanovic (2018). They differ in the number of classes and the methods used to construct income classes.

We incorporated all social table studies available in the public domain, but acknowledge that we may not have an exhaustive list. Crucially, our data set includes information on all the social tables used by Milanovic (2018). Table C.2 in Appendix C lists the full set of sources we used to compile our dataset. We also add social table information on Africa and South America, regions that were mostly or completely absent from Milanovic’s earlier work. Including more tables increases the statistical power of the analysis and improves geographic coverage, giving a more nuanced picture of inequality and social tables inequality studies across the globe than was previously possible.

Our dataset includes information on the author of each social table, the year of publication, the region and country being studied, the year for which the inequality metric is calculated, the population size, the surface area of the country examined, the Gini coefficient calculated from each social table, the number of social classes, the proportion of the population represented by the top and bottom classes, and GDP per capita at 1990 prices from the Maddison project (Bolt and Zanden 2013)⁵. We used these variables to discern systematic patterns that were related to the country context and the social table attributes used to measure inequality. A unique feature of this data set is that we include studies which calculate social tables for various countries over several periods, such as (Bolt and Hillbom 2016; Aboagye and Bolt 2021; Haas 2017). This enables us to ascertain whether adding inequality trends from the global South to social table analysis changes our perspectives on inequality and development broadly, and on the social tables method more specifically. Our larger sample allows us to compare social tables constructed as early as 1962 with those constructed more recently and for different global regions. This is particularly important because data quality can vary substantially across regions.

For studies that report population shares instead of total populations, such as (FitzGerald 2008), we use the latest population figures from the Maddison project database (Bolt and Zanden 2013). Some social tables, such as that for Java in 1880, report the share of households in each class instead of the share of people. These instances are rare and we argue that they should not systematically bias results. Like (Milanovic 2018), we use today’s country surface areas where historical information is not available.

III Analysis

III.I Repeated Case Studies

As a starting point, figure 4 presents country-specific examples that provide motivation for our hypothesis: that the size of the bottom class matters in a case with a small population. It shows the results of studies of four societies for each of which we have collected two sets of matching inequality estimates that should – apart from data limitations – be identical. They share the same historical reference year, but the estimates

⁵Although GDP per capita data at 2011 prices is available, we use the 1990 PPP estimates to keep them comparable with figures drawn from Milanovic (2018)

were released to the public at different times. Authors used alternative assumptions to construct the tables. In some cases the same authors generated the two sets of estimates using similar archival sources – typically one estimate is lifted from an early working paper version, while the alternative is an update to the analysis. Some of the “second” versions have been peer reviewed for publication. In other cases two different authors studied the same society and period, but constructed their tables independently of each other.

For example, in the top left panel (a), we compare two peer reviewed inequality estimates for the same population from Chile in the same historical reference period of 1900. The first was published by FitzGerald (2008), and the second was absorbed into the meta-analysis of Milanovic (2018). The 2008 estimate uses only four classes, with the bottom class contributing as much as 41% to the population distribution and the top class 8%. The resulting Gini coefficient is moderately low at 0.365, raising concerns that within-group variation has been neglected. The 2018 estimate includes more variation by using ten classes, with a much smaller bottom class (2.6%). But the top class accounts for as much as 36% of the population. Now the Gini estimate climbs to 0.45. This example supports the common understanding that adding more classes increases between-group variation and reduces the risk of missing differences within these groups. It supports our hypothesis that using large bottom classes results in downward bias of inequality estimates – by as much as 10 percentage points. However, the comparison is inconsistent with our expectations that large top classes could also bias social table inequality estimates downward. Be that as it may, this discrepancy is congruent with Modalsli (2015) findings that the size of the top class is not important for the robustness of inequality estimates. Further analysis will be required to investigate this issue further.

The top right panel (b) compares a pair of working papers (dashed lines and empty dots) and a pair of published estimates (solid lines and solid dots) for the same years from 1891 to 1960 in Ghana (Aboagye and Bolt 2021). The estimates cover a range of periods, using a comparable set of archival sources in each period. This case illustrates the effect of differences in the size of the bottom classes: they were smaller in the earlier versions of the research. By contrast, the number of classes remained constant at around 16 to 17 in all periods and versions of the analysis. The size of the top class was consistently small in both versions (<1%). This case therefore more readily isolates a negative association between the size of the bottom class and the inequality estimates. The earlier Gini estimates were mostly higher when the bottom class was smaller. These observations make a case for our hypothesis that even when researchers choose a large number of classes, a large bottom class produces sensitive results.

The bottom panels, using data from (c) Senegal and (d) the Ivory Coast (Alfani and Tadei 2017), further discredit the view that the number of classes is the most important determinant of bias in estimating inequality with social tables. This conclusion is in line with multi-country evidence presented by Modalsli (2015) and Milanovic (2018). Alfani and Tadei’s (2017) earlier versions of the analyses relied on only about 10 classes. In later versions of the same paper they increased the number to 40. Yet they found, especially in the case of the Ivory Coast, that the Gini coefficients were insensitive to this large increase. The size of the bottom class fluctuated, both over the period analysed and across the versions of the analysis, but the inequality estimates nevertheless remained relatively stable. However, the gap across approaches grew larger from the 1950s when the Senegalese data did not allow the researchers to break down the bottom class significantly.

Conclusions from the repeated case studies point to a possible systematic relationship between social table attributes and inequality estimates, though the nature of the bias caused by these relationships is confounded by context-specific factors. This calls for a more systematic cross-country and -context specific investigation of the general trends in social table analysis.

III.II Evolution of social table attributes

Before we look at systematic bias in inequality estimates, we gauge researchers' responsiveness to concerns that inequality estimates are sensitive to social tables attributes. Figure 5 tracks population shares in top and bottom classes and the number of classes over two trajectories of interest. Firstly, we follow social table attributes according to the year in which the table was constructed or published (depending on which date is available) – this view primarily captures the evolution of author behaviour, from the earliest papers to later studies. Secondly, we present the figures according to the historical period analysed – this shifts our perspective to understanding changes related to economic development and data quality. Comparing the two trajectories helps us distinguish whether researcher awareness or the analysis of richer sources from later periods is the primary determinant of the evolution of the type of social tables being produced.

Panels (a) and (b) of figure 5 show that average population sizes have remained stable across publication years and historical years. Because other attributes – such as number and size of classes – evolve more significantly, similarly sized distributions are being pieced together in systematically different ways across the periods analysed, and by researchers working in different times. The broader context has also changed. Panel (c) of figure 5 shows that more recent studies have increasingly looked at less equal societies. More recent research has therefore focused on compiling historical evidence for places with high inequality. Panel (d) shows that the lowest levels of inequality were detected in countries that existed around the turn of the 21st century. Researchers' attention shifted back and forth between societies with higher and lower GDP, but the most recent historical settings (approaching the year 2000) have the highest GDP and the lowest Gini coefficients.

Panel (a) shows that earlier analyses deployed more – and not fewer – classes. This is closely reflected in the use of fewer classes to re-construct distributions of more recent historical periods. Social tables research has therefore shifted to societies with limited data. These trajectories suggest that data availability is a greater constraint than researcher awareness in constructing social tables. The evidence suggests that even the oldest studies have documented as many classes as the data has allowed, but that newer studies have tried to unlock evidence for contexts with higher data demands. Both trajectories show that more social tables with smaller bottom classes have emerged from sources in recent years. Both data availability and researcher awareness could have produced this trend. By contrast, panel (a) shows that, since the turn of the century, researchers have consistently paid attention to using smaller top classes. There are exceptions in studies of countries at a more recent date whose archival sources have limited information on somewhat larger top classes, as can be seen in panel (b). By this count, it is again clear that researcher awareness is not the main constraint to constructing classes that reflect a population accurately. Social tables published since 2000 seem to have focused on building distributions with smaller top classes; however, this approach has not necessarily produced systematically different Gini estimates – as Modalsli (2015) has concluded.

III.III Associations between social table outcomes and attributes

III.III.1 Bivariate analysis

The structure of the social tables emerging in the literature has evolved – mainly because of changing source quality for societies in recent historical periods, but also possibly because researchers' awareness of potential biases has increased. Changes in measured inequality may be partially attributed to the changes in the way social tables have been constructed. To explore this possibility, figure 6 which we consider to be representative of the earliest literature on social tables, and which focused mainly on the pre-industrial era. Each panel shows three samples. The first is limited to 28 studies analysed by Milanovic, Lindert, and

Williamson (2011), which we consider to be representative of the earliest literature on social tables, and which focused mainly on the pre-industrial era. The second uses an expanded sample analysed by Milanovic (2018) and represents the evolution of the research landscape that includes newer methods and more studies from the period of industrialization. The third sample is our full data set of 108 studies. The additions to this data set are a mix of older studies (not previously included by Milanovic and his colleagues) and new studies that have expanded the geographic scope to developing areas in Africa (Bolt and Hillbom 2016; Aboagye and Bolt 2021; Haas 2022) and Latin America (Bértola et al. 2010; FitzGerald 2008; Bertlola, Gelman, and Santilli 2015; Castañeda Garza and Bengtsson 2020). Including these societies in the analysis is arguably important because they have high levels of inequality today. As mentioned above, newer social tables have analysed countries with higher inequality and somewhat lower economic development. Because these countries also have poorer data quality, the tables have fewer classes, and consequently larger bottom classes. Table B.1 in appendix B reproduces Milanovic (2018)’s regression specifications as far as possible with similar samples as in figure 6. This supplementary analysis confirms that our sample is representative of earlier results similar to Milanovic (2018). However, we add datapoints from new studies for analysis, which increases the statistical power and improves the ability to detect more complex relationships.⁶ We refer to the data used by Milanovic, Lindert, and Williamson (2011) as the 2011 sample, the data used by Milanovic (2018) and our data set as the 2022 sample.

Bivariate associations in figure 6 (a) confirm the negative parametric relationship between population and inequality found by Milanovic (2018). But, as our 2022 sample shows, this relationship is non-linear. Similarly, panel (b) shows that introducing more variation across time and space yields the well-known inverted U-shaped Kuznets curve (Kuznets 1955). A first unsurprising lesson is that cross-country analysis requires sufficient statistical power to tease out relationships. However, a more pertinent second lesson sounds a warning. Our analysis in figure 5 shows that the inverted U may very well be a function of observing a larger set of new studies that analyse more recent historical periods and countries from the global South. These studies correspond to periods of higher growth globally (and not specifically in these generally poorer regions), but also to social table attributes. Whether the Kuznets curve emerges because of statistical biases stemming from the use of social tables, or whether this reflects reality, remains to be established.

Panel (c) of figure 6 prompts us to consider the latter possibility more seriously. Some studies in our expanded sample have high Gini coefficients. They also happen to have been constructed from a growing number of classes. This is consistent with the view that measured inequality is higher when more classes are used. The observed “Kuznets” relationship could, therefore, be merely an artefact of changing estimation conditions.

Panel (d) confirms Modalsli’s (2015) previous findings that social table Gini estimates are insensitive to a large top class. Panel (e) shows that the proportion of the population in the top class also does not vary significantly with population density. In contrast, and central to our argument, panel (f) shows that a large – and especially an extremely large – bottom class produces lower Gini estimates. Both older and newer studies fit this pattern. Panel (g) suggests that using a large bottom class is closely related to having only a few classes in the social table. These findings suggest that while having few classes may not directly affect inequality estimates, it introduces an indirect mechanism that may do so: the presence of a large bottom

⁶For the sake of comparison, following Milanovic (2018) we do not log population density and GDP. We use 1990 PPP GDP estimates from Madison’s data set to be able to compare our results with earlier analysis. Our observation count is smaller than Milanovic’s for a number of reasons. Firstly, we could not locate all the estimates in his original data. Secondly, we had missing data on land area for some of the territories whose boundaries we cannot recover; this meant we could not calculate population densities for these observations. Thirdly, we excluded some of Milanovic’s inequality estimates constructed using historical micro data from some of our specifications. These particular estimates do not qualify as “social tables”.

class. The non-significant relationship between top and bottom classes that we see in (i) reinforces our findings from panel (d) that the large bottom class problem is not mirrored at the top of the distribution. Furthermore, there is also a definite negative relationship between a very high population density and a large bottom class share, as panel (h) shows. This warrants specific investigation of societies with smaller populations, to establish whether large bottom class shares are more serious for estimation bias in these contexts.

III.III.2 Multivariate analysis

Table 1 draws our findings together in a multivariate framework. We estimate Ordinary Least Squares (OLS) equations on our full sample as well as the limited sample of Milanovic (2018). We adapt this functional form to test for the Kuznets relationship, while controlling for social table attributes. Our interest is twofold: firstly, we are interested in observing any systematic biases on social table attributes; secondly, we assess whether the Kuznets relationship observed in figure 6 is completely attributable to changes in the way social tables have been constructed over time, or whether it is a statistical reality.

$$\begin{aligned}
 gini_{ict} = & \beta_0 + \beta_1 \log(GDP \text{ per capita})_{ict} + \beta_2 \log(GDP \text{ pc})_{ict}^2 + \beta_3 \log(\text{population density})_{ict} \\
 & + \beta_4 \log(\% \text{ in bottom class})_{ict} + \beta_5 \log(\% \text{ in top class})_{ict} + \beta_6 \log(\text{number of classes}) \\
 & + \gamma_c + \varepsilon_{ict}
 \end{aligned}$$

where i indexes the region, c the continent, t the period and γ_c are continent fixed effects.

We adapt this specification in two ways to explore the hypothesis in greater depth. Firstly, we introduce interactions of the social attributes with dummy variables, $I(GDPpc > \$4000)$ and $I(Population > 2million)$, to test our hypothesis that fewer, larger classes introduce significant biases in characterizing distributions only in small populations or societies with high economic development. Secondly, we assess whether older publications are more or less prone to such bias by estimating rolling regressions. We estimate all these regressions using the entire sample to preserve statistical power. However, in each year of interest we re-weight data points by their proximity to the reference period: $w_{it} = \frac{1}{|year_t - reference_t|}$. In other words, the further an observation is in time from a particular year, the lower the weight it carries in the comparison.

The first three columns of table 1 use the 2018 sample. Regardless of whether or not we include micro data estimates in our sample, we detect no significant relationships between the covariates and the social table inequality estimates. However, the null results may be solely a function of a lack of statistical power. We therefore expand our analysis to incorporate our full 2022 sample in column (4). This simplest specification confirms our findings from the bivariate analysis: there is a cross-country historical Kuznets curve with no relationship between population density and inequality. Column (5) excludes observations derived from micro data, and the findings remain robust. The Kuznets curve is therefore not produced by including systematically different micro data inequality estimates in the analysis. Column (6) takes the limited 2018 sample of social tables inequality estimates forward and introduces controls for social table attributes. The inverted U-shaped relationship between GDP and inequality remains significant and stable, though the social table attributes do not predict inequality and also do not artificially produce a Kuznets relationship.

Given the context-specific problems we identified in figure 4, we continue to explore whether there are cases that deviate from the statistical patterns estimated in this specification. Figure 7 plots leverage statistics against a range of covariates. While leverage is mostly unrelated to population size, we observe a distinct negative relationship between countries with populations of less than two million. Similarly, leverage

has a positive association with GDP per capita, but only beyond the threshold of US\$4000 per person. The changes at these thresholds indicate that the social table attributes have non-linear and heterogeneous relationships with the Gini coefficient. We explore these possibilities by adapting specifications (7) and (8) to interact social table attributes with indicators for these thresholds. While we observe no significant heterogeneity by economic development, we find that using larger bottom classes produces significantly lower Gini estimates only when total populations are small (below two million).

We also find non-linear relationships between social table attributes and leverage. We adapt our models for non-linearity in these variables, but they do not significantly and materially alter our conclusions. Nevertheless, we observe distinctly high leverage for observations with extremely low bottom class shares. This supports our findings that population size and bottom class shares interact to generate significantly different Gini estimates.

In figure 8 we show that the associations between selected social table attributes – the number of classes and the proportion of the population in the top class – and inequality are stable over time; they are non-significant regardless of the publication date. This once again confirms Modalsli’s (2015) conclusion that the size of the top class is of little concern for the construction of social tables. Nevertheless, more recently published studies are providing the statistical power that reveals more clearly the combined role of small populations and large bottom classes in creating a downward bias in social tables inequality estimation. We interpret this in light of equations 2 and 3, which show that even small variations in income within the bottom classes can have a large effect on the total population Gini coefficient. When bottom shares are large, they amplify the effect of small biases. Our findings show that the effect holds only in small populations of less than 2 million inhabitants. These contexts are associated with having few social classes, low sectoral diversity and larger bottom classes. While our evidence suggests that small bottom classes are not a serious concern in most estimation contexts, they should be taken seriously when small populations have large subsistence or informal sectors. Our derivations show that subdividing large bottom classes so as to distinguish small differences can produce better estimates of inequality.

IV Conclusion

Economic historians have made large strides in the past two decades in calculating more accurate estimates of inequality over the (very) long run. New research has produced inequality estimates ranging from pre- to post-industrial periods and covering diverse geographic areas. However, there is more to be done. Scant or poor quality micro-level data limits the periods and places for which we can make global comparisons. This is particularly true of the global South, where inequality estimates are sometimes hard to come by, even for the present day. Researchers are exploring the possibility of using social tables to remedy this deficiency. This approach is suitable for reconstructing distributions from a range of sometimes unconnected archival sources, which – in contrast to micro records – are possible to assemble retrospectively. Researchers’ enthusiasm for broadening the evidence base with social tables continues to grow, but they continue to be uneasy about relying on simplifying assumptions which raise questions about their credibility. Using only a few representative groups to compile complete income distributions has its limitations, chief among which is the default assumption that within-group inequality is zero. Adding more classes can mitigate this problem, but early research suggested that having only a few classes may not be as serious a shortcoming as was expected (Milanovic, Lindert, and Williamson 2007). A large population share in the top class also seems to be surprisingly unproblematic, though overlap across groups is more serious (Modalsli 2015).

Previous attempts to stress-test social table assumptions have used only a small set of estimates from a few studies. The present paper extends this line of inquiry in two ways. Firstly, we surveyed a broad

evidence base, consisting of 108 social table estimates from pre-industrial and post-industrial societies; a much larger number than the 41 analysed in Milanovic’s (2018) comparative study. The larger database increases the statistical power for teasing out relationships and improves geographic coverage to include more recent evidence from the global South. Secondly, we turned our attention to the role of large bottom classes. This feature of social tables has been neglected in favour of concerns with large top classes, partly because of their demonstrated relevance in the recent literature (Piketty, Saez, and Zucman 2022). We show there are strong reasons to believe that a large bottom class may introduce a meaningful downward bias to inequality estimates, even in cases with relatively low variation in the bottom tail. Cases of this kind are relevant for research on societies with large homogeneous subsistence sectors.

We first looked at the effects of different ways of measuring inequality in the same geographic settings and periods, using some repeated case studies. These showed that inequality estimates are sensitive to how researchers construct their social tables. The estimates remained stable when more classes were added, using the same archival sources. However, when the tables were reconfigured and the size of the bottom class increased, the inequality estimates showed a downward bias. The way social tables are structured is therefore important for producing credible inequality estimates. The evolution of social table structures over time suggests that researchers have become aware of the possible problems. However, their choices are constrained by the availability and quality of archival data. Contexts with data limitations have obliged them to build simpler tables.

Deeper investigation using bivariate and multivariate statistics did not show that introducing more classes directly reduced the downward bias introduced by omitting within-class variation. We found no evidence to suggest that researchers have used excessively large top classes. This is surprising, given the recent attention to distinguishing the “very rich” from the “super rich” in studying distribution, but it is to be expected since archives document the lives of the elite better than those of the poor, particularly in colonial settings (Gordon 2018; Fourie and Green 2015). It is thus easier to distinguish variation at the top of the income distribution than at the bottom. It is especially difficult to distinguish between similarly poor groups in societies with small populations, and we found that lumping them together can create a downward bias in estimating inequality. Typically, this problem could arise in economies that are sectorally homogeneous, where it is difficult to distinguish even broad differences in socioeconomic status in the subsistence sector, let alone fine differences. We stress the importance of trying to establish income differences in bottom tails, despite the difficulties presented by insufficiently detailed archival accounts of the socioeconomic conditions of the poor, and particularly in the case of indentured and enslaved people who did not earn incomes (Fourie and Fintel 2010).

Given that researchers are increasingly using the social tables method to expand global inequality estimates to include data scarce contexts, we need to find creative and rigorous ways to overcome the limits of archival sources. These could include developing better statistical methods, exploring new archival sources, and agreeing on conventions for choosing the size of social table classes to minimize bias, while ensuring that estimates from various settings are harmonized. These proposals are less onerous for producing credible comparisons within countries or regions, but producing comparable estimates that have global coverage over the long run will require significant agreement and cohesion. This paper has built on previous contributions that have examined the rigour of estimating inequality with grouped data (Milanovic, Lindert, and Williamson 2007; Modalsli 2015; Van Ourti and Clarke 2011; Shorrocks and Wan 2008; Dagum 1997; Kakwani 1980; Gastwirth 1972; Cowell 1991). In re-assessing unresolved questions, we have focused on variation in the bottom tail, a part of the puzzle that may matter in contexts where this type of research is taking off. We also make a call for greater cohesion and agreement in producing comparable estimates that can give a global

view of historical inequality. Current initiatives to co-ordinate such efforts are a step in the right direction.

Figure 1: Regional evolution of social table studies, by date table was created

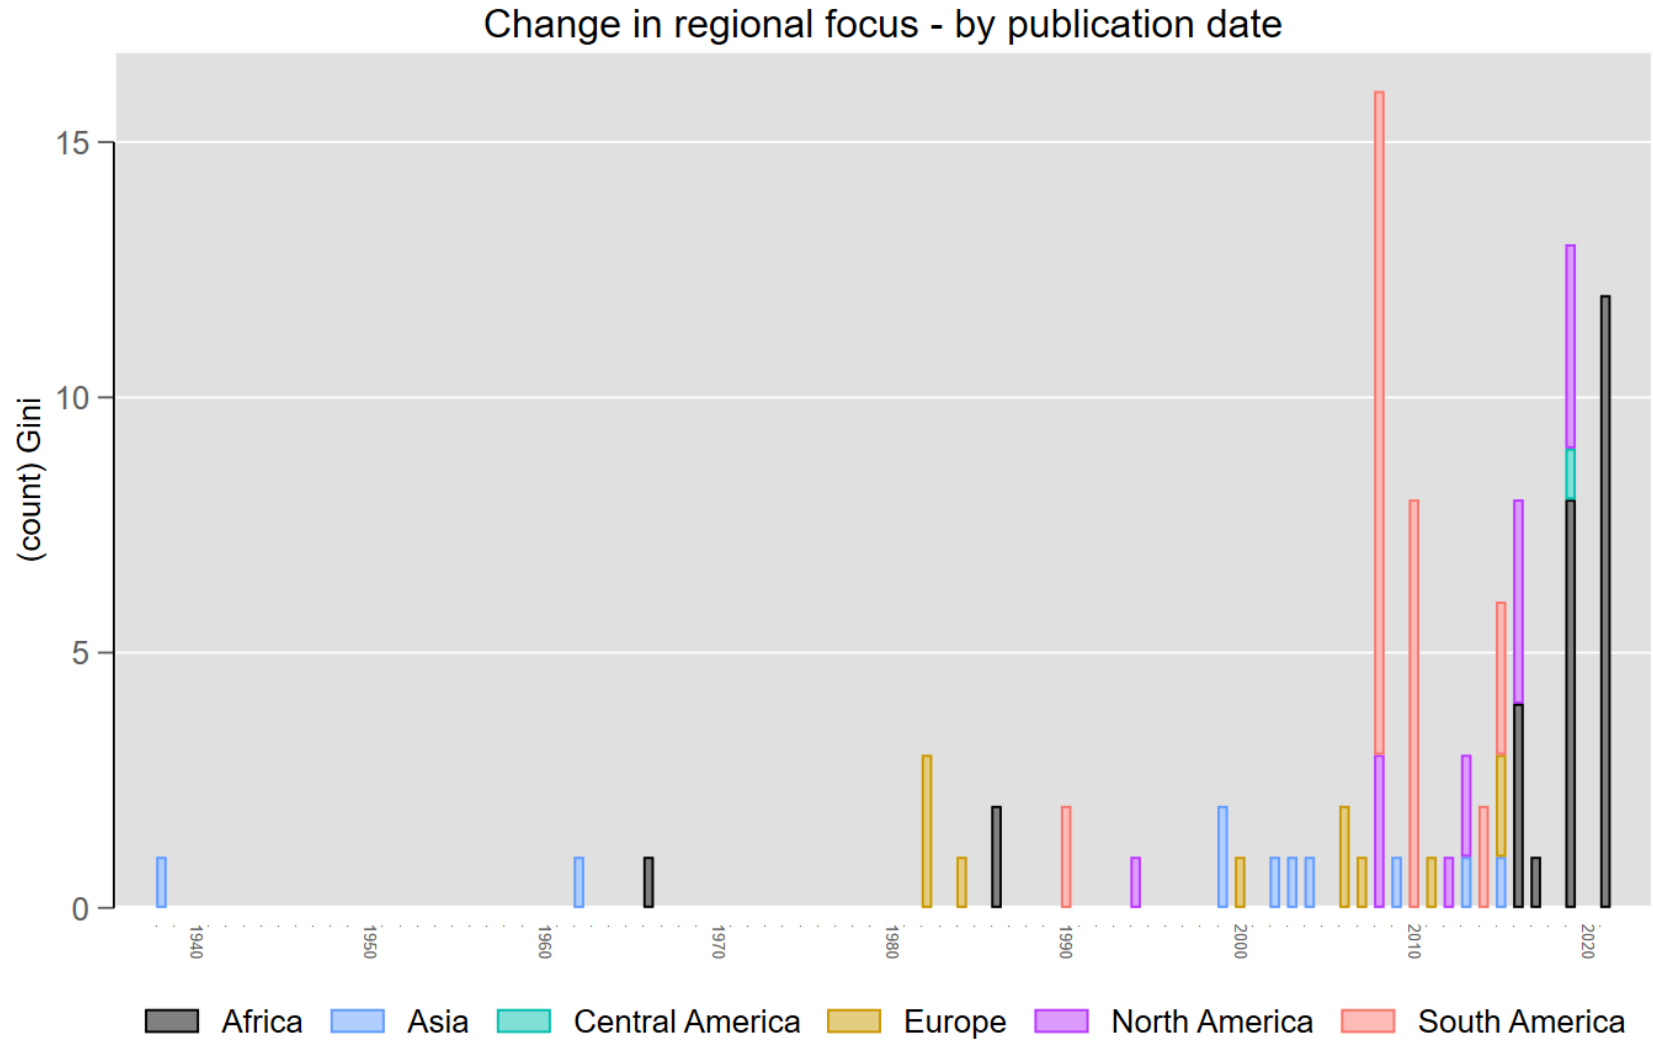


Figure 2: Number of social table studies for pre-1800 societies

Number of Social Table Studies for pre-1800 societies




Figure 3: Number of social table studies for post-1800 societies

Number of Social Table Studies for post-1800 societies

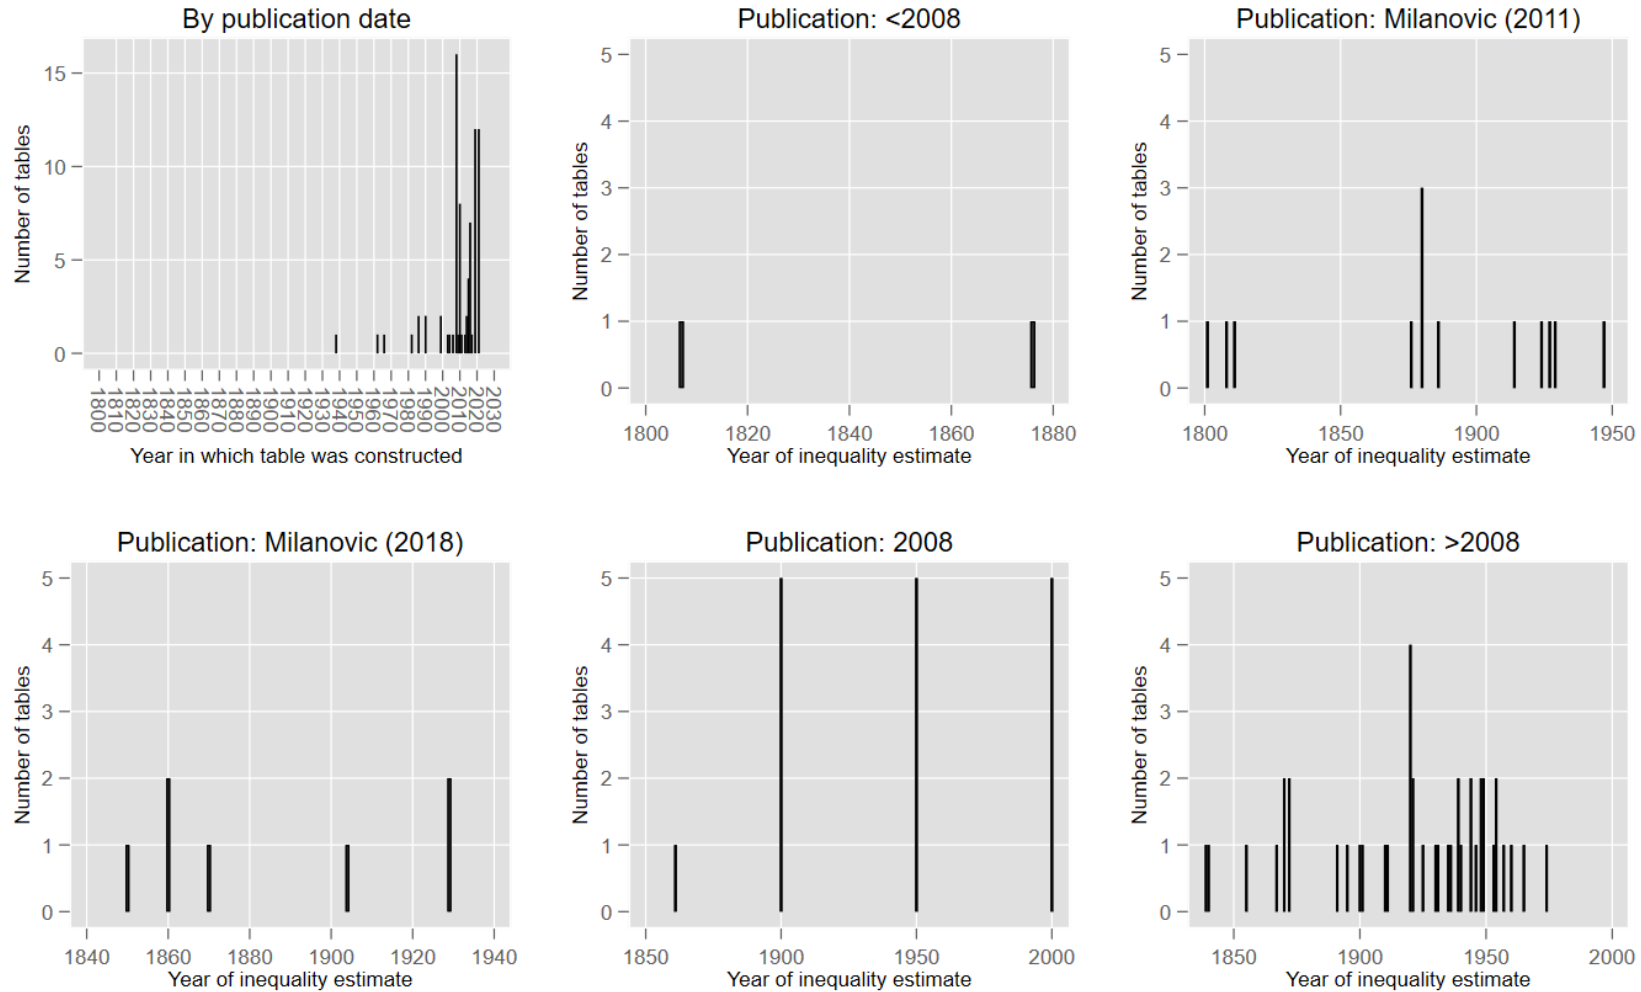


Figure 4: Gini estimates from same contexts but with different social table attributes

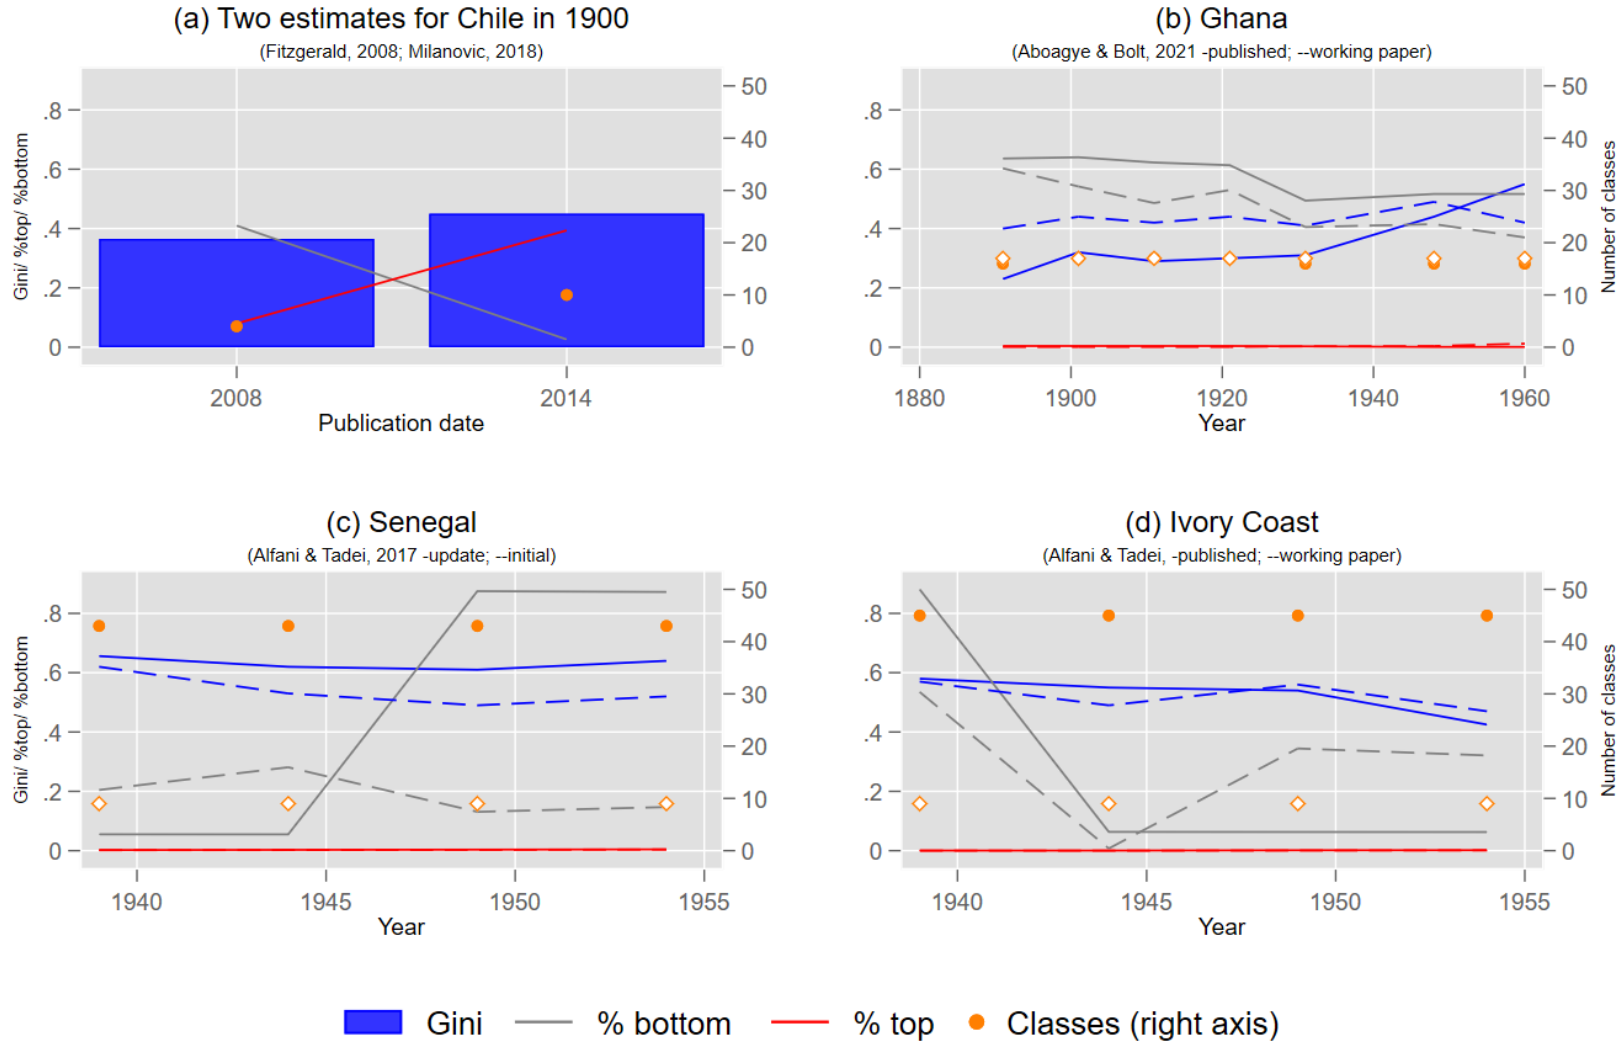


Figure 5: Evolution of social table attributes

Characteristics of Social Tables by publication date or reference year


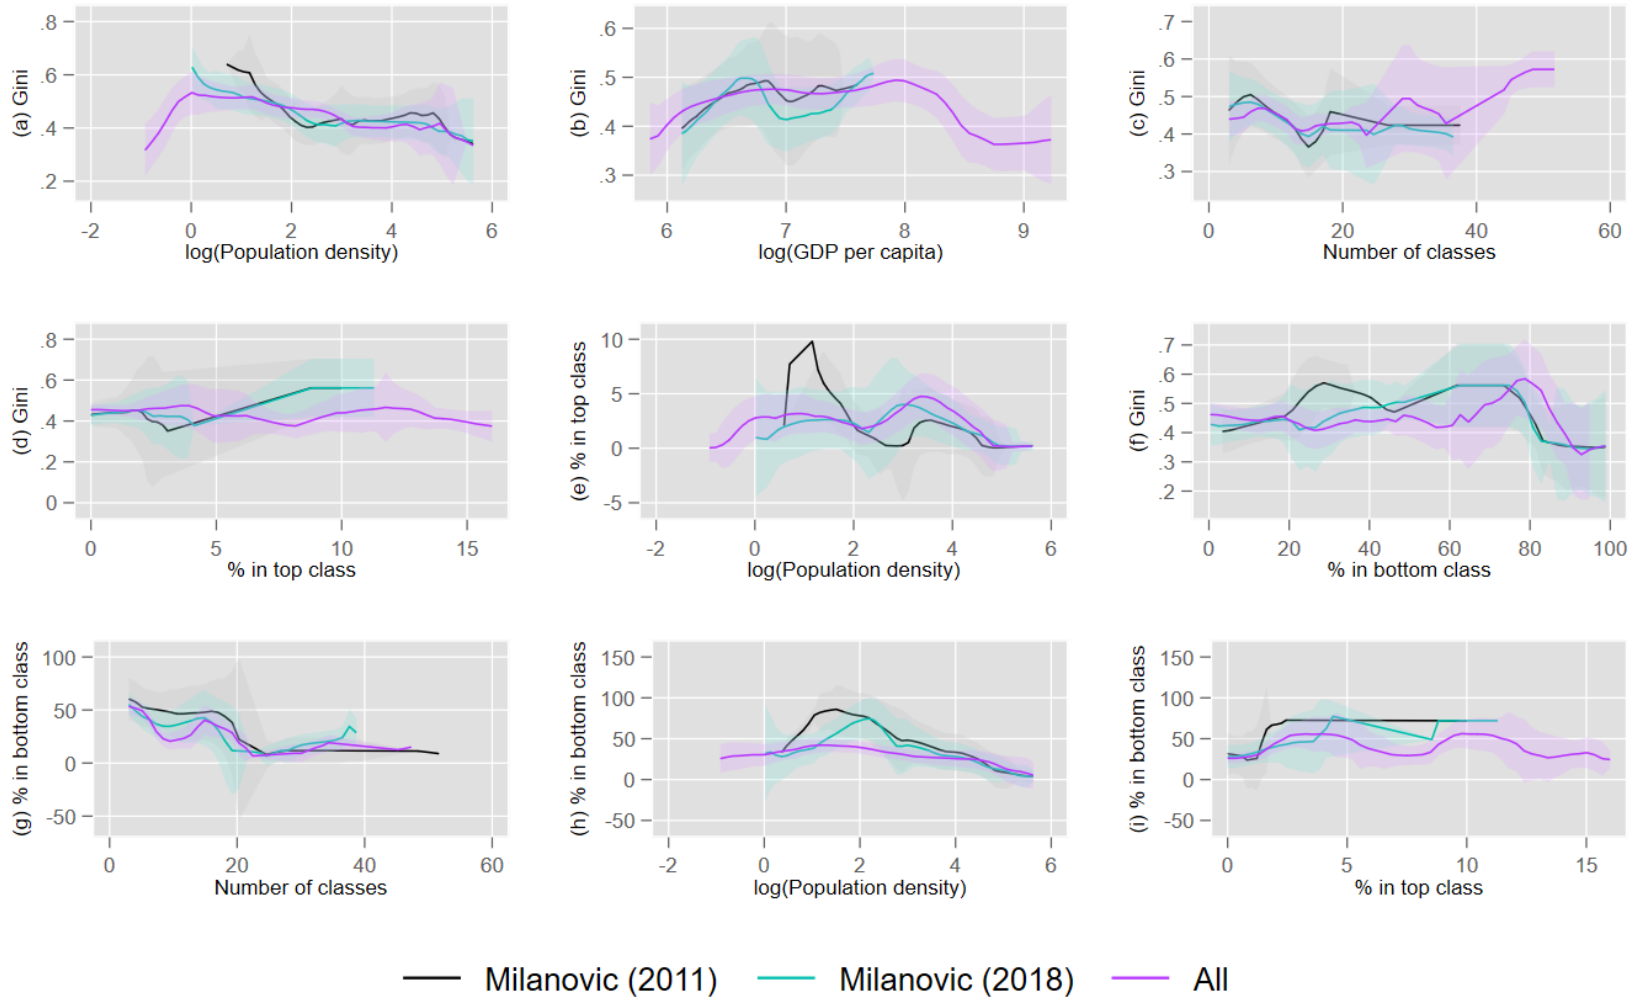


Figure 6: Association between social table outcomes and attributes

Associations between social table outcomes and attributes



Sample	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Milanovic (2018)			All				
log(Population/km ²)	-0.001 (0.029)	-0.021 (0.040)	0.037 (0.053)	-0.016 (0.012)	-0.018 (0.012)	-0.020 (0.014)	-0.017 (0.014)	
log(GDP pc)	0.530 (1.369)	1.100 (1.565)	0.045 (1.773)	0.807*** (0.236)	0.723*** (0.235)	0.651** (0.269)		0.579** (0.279)
log(GDP pc) ²	-0.036 (0.098)	-0.079 (0.113)	-0.010 (0.126)	-0.056*** (0.016)	-0.050*** (0.016)	-0.045** (0.018)		-0.041** (0.019)
log(% in top class)			0.012 (0.009)			-0.002 (0.006)	-0.003 (0.006)	-0.012 (0.019)
log(% in bottom class)			0.034 (0.030)			0.002 (0.012)	0.001 (0.012)	-0.082** (0.037)
log(classes)			0.098 (0.060)			0.003 (0.019)	0.007 (0.019)	-0.052 (0.060)
I(GDPpc>4000)							-0.120 (0.592)	
I(GDPpc>4000) × log(% in top class)							0.010 (0.174)	
I(GDPpc>4000) × log(% in bottom class)							-0.026 (0.205)	
I(Population > 2million)								0.017 (0.106)
I(Population > 2million) × log(% in top class)								0.011 (0.020)
I(Population > 2million) × log(% in bottom class)								0.092** (0.039)
I(Population > 2million) × log(classes)								0.058 (0.061)
Constant	-1.437 (4.724)	-3.301 (5.385)	0.224 (6.114)	-2.393*** (0.875)	-2.066** (0.875)	-1.832* (0.986)	0.463*** (0.054)	-1.579 (1.047)
Include micro estimates	Y	N	N	Y	N	N	N	N
Region FE	Y	Y	Y	Y	Y	Y	Y	Y
N	33	26	25	80	73	64	64	69

Standard errors in parentheses
* p<0.1, ** p<0.05, *** p<0.01

Table 1: Results

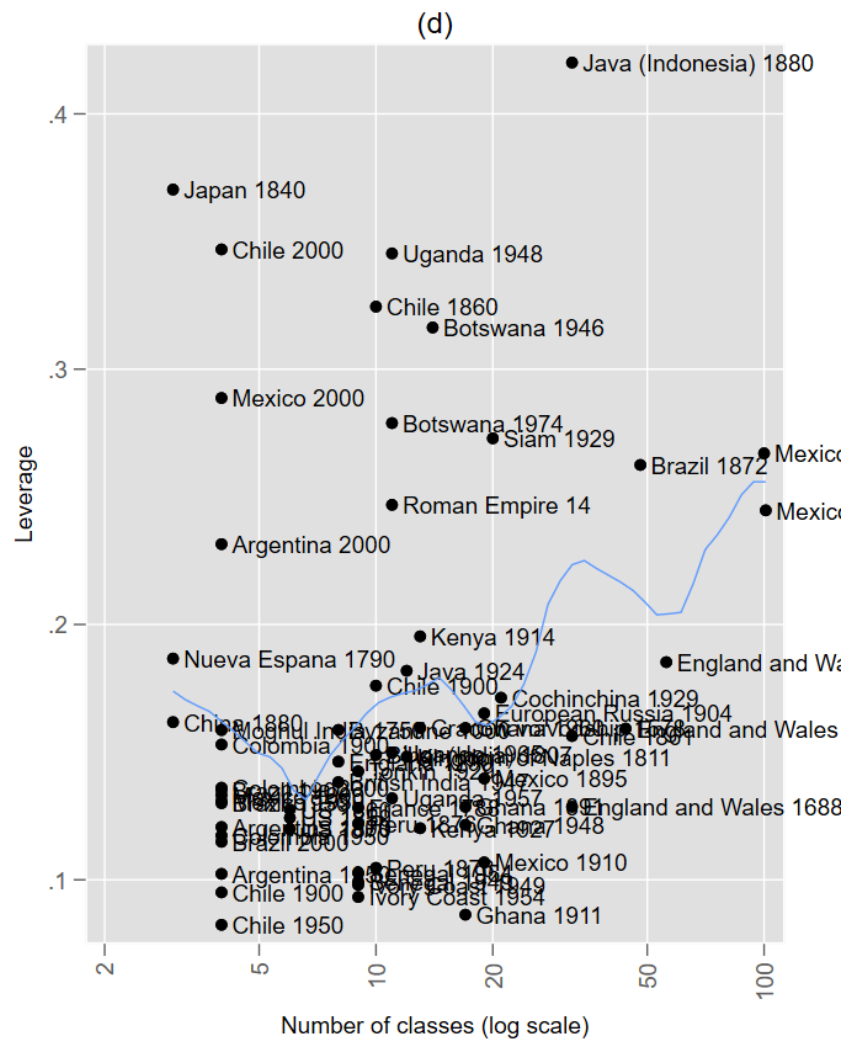
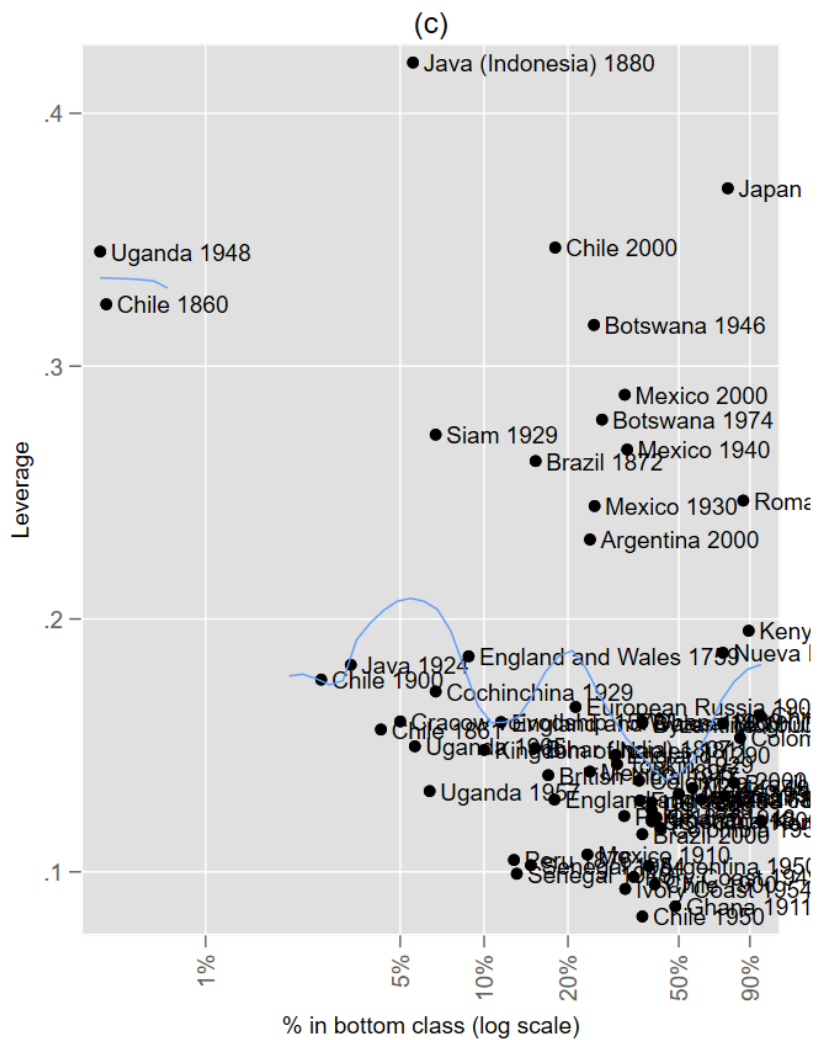


Figure 7 cont.: Association between leverage statistics from table 1, column (6) and other covariates

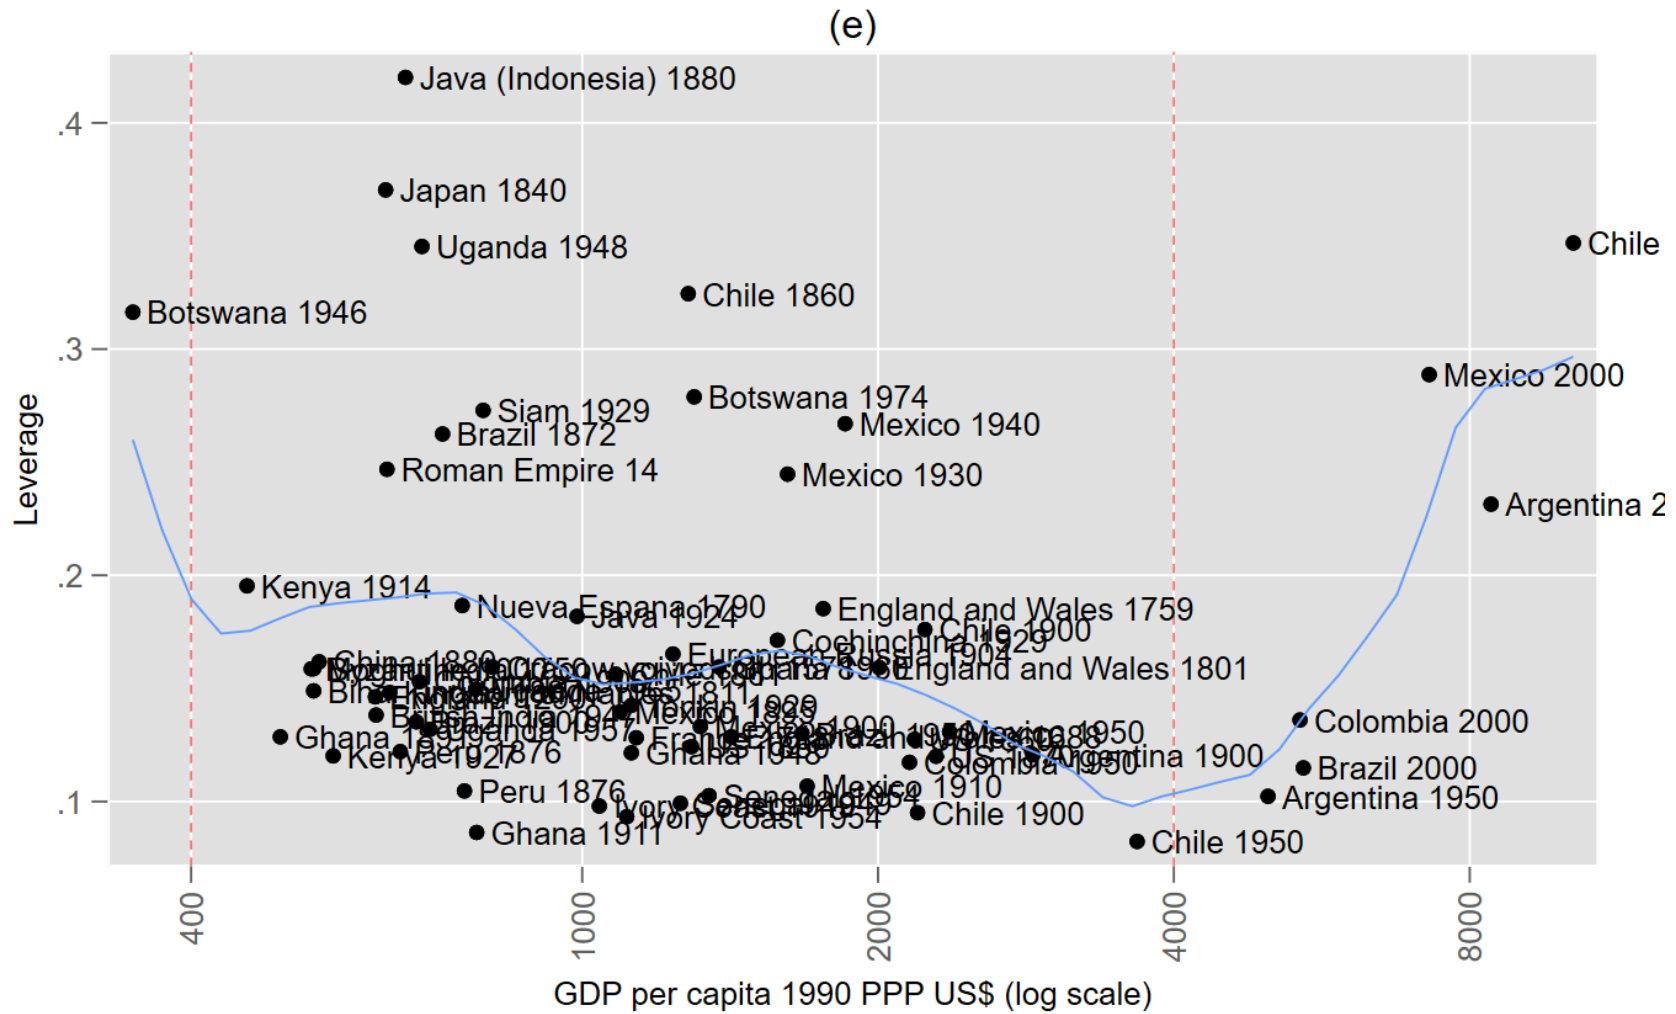


Figure 7 cont.: Association between leverage statistics from table 1, column (6) and other covariates

Rolling regression coefficients on Gini: by Publication Date

with 95% confidence interval

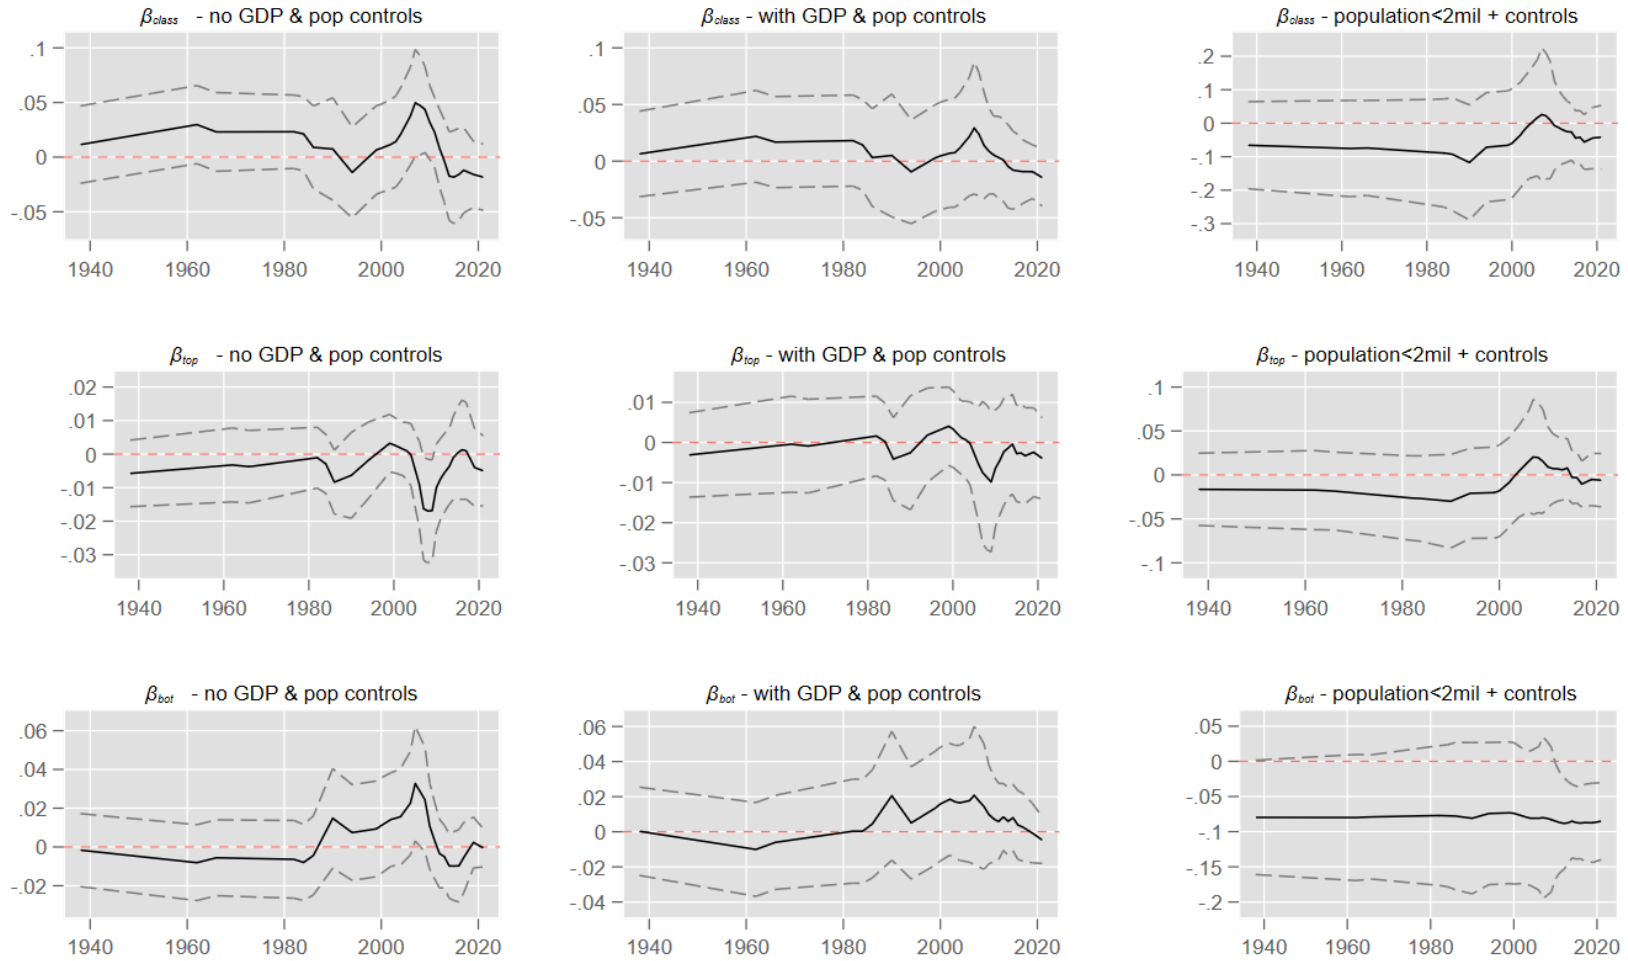


Figure 8: Rolling regression coefficients of Gini on social table attributes

References

- Aboagye, Prince Young and Jutta Bolt (2021). “Long-term trends in income inequality: winners and losers of economic change in Ghana, 1891–1960”. In: *Explorations in Economic History* 82, p. 101405.
- Aghevli, Bijan B and Farhad Mehran (1981). “Optimal grouping of income distribution data”. In: *Journal of the American Statistical Association* 76.373, pp. 22–26.
- Alfani, Guido (2021). “Economic Inequality in Preindustrial Times: Europe and Beyond”. eng. In: *Journal of economic literature* 59.1, pp. 3–44.
- Alfani, Guido, Federico Tadei, et al. (2017). “Income inequality in colonial Africa: Building social tables for pre-independence Central African Republic, Ivory Coast, and Senegal”. In: *African Economic History Working Paper* 33, p. 2017.
- Alvaredo, Facundo (2011). “A note on the relationship between top income shares and the Gini coefficient.” In: *Economics Letters* 110.3, pp. 274–277.
- Alvaredo, Facundo, Anthony B Atkinson, Thomas Piketty, and Emmanuel Saez (2013). “The Top 1 Percent in International and Historical Perspective”. In: *Journal of Economic Perspectives* 27.3, pp. 3–20.
- Atkinson, Anthony B (2007). “Measuring Top Incomes: Methodological Issues”. In: *Top Incomes over the Twentieth Century: A Contrast Between Continental European and English-Speaking Countries*. Ed. by Thomas Piketty and Anthony Atkinson. Oxford University Press.
- Atkinson, Anthony B, Thomas Piketty, and Emmanuel Saez (2011). “Top incomes in the long run of history”. In: *Journal of economic literature* 49.1, pp. 3–71.
- Berry, Robert (1993). “Civil Liberties Constraints on Tribal Sovereignty After the Indian Civil Rights Act of 1968”. In: *JL & Pol’y* 1, p. 1.
- Bertola, Luis, Jorge Daniel Gelman, and Daniel Victor Santilli (2015). “Income distribution in rural Buenos Aires, 1839-1867”. In:
- Bértola, Luis, Cecilia Castelnovo, Javier Rodríguez, and Henry Willebald (2010). “Between the colonial heritage and the first globalization boom: on income inequality in the Southern Cone”. In: *Revista de Historia Económica* 28.2, p. 307.
- Bolt, Jutta and Ellen Hillbom (2016). “Long-term trends in economic inequality: lessons from colonial Botswana, 1921–74”. In: *The Economic History Review* 69.4, pp. 1255–1284.
- Bolt, Jutta and Jan Luiten van Zanden (2013). “The maddison project”. In: *Retrieved October 14*, p. 2019.
- Booth, Anne (1988). “Living standards and the distribution of income in colonial Indonesia: A review of the evidence”. In: *Journal of Southeast Asian Studies*, pp. 310–334.
- Castañeda Garza, Diego, Erik Bengtsson, et al. (2020). *Income Inequality in Mexico 1895-1940: Industrialization, Revolution, Institutions*. Tech. rep. Lund University, Department of Economic History.
- Cowell, Frank A (1991). “Grouping bounds for inequality measures under alternative informational assumptions”. In: *Journal of Econometrics* 48.1-2, pp. 1–14.

- Dagum, Camilo (1997). “A New Approach to the Decomposition of the Gini Income Inequality Ratio.” In: *Empirical Economics* 22.4, pp. 515–531.
- FitzGerald, Valpy (2008). “Economic development and fluctuations in earnings inequality in the very long run: The evidence from Latin America 1900–2000”. In: *Journal of International Development: The Journal of the Development Studies Association* 20.8, pp. 1028–1048.
- Fourie, Johan and Dieter von Fintel (2010). “The dynamics of inequality in a newly settled, pre-industrial society: the case of the Cape Colony”. en. In: *Cliometrica* 4.3, pp. 229–267. DOI: 10.1007/s11698-009-0044-1.
- (2011). “A History with Evidence: Income Inequality in the Dutch Cape Colony”. In: *Economic History of Developing Regions* 26.1. DOI: 10.1080/20780389.2011.582990.
- Fourie, Johan and Erik Green (2015). “THE MISSING PEOPLE: ACCOUNTING FOR THE PRODUCTIVITY OF INDIGENOUS POPULATIONS IN CAPE COLONIAL HISTORY”. In: *The Journal of African History* 56.2, pp. 195–215.
- Galli, Stefania, Dimitrios Theodoridis, and Klas Rönnbäck (2022). “Economic inequality in Latin America and Africa, 1650 to 1950: Can a comparison of historical trajectories help to understand underdevelopment?” In: *Economic History of Developing Regions*, pp. 1–24.
- Gastwirth, Joseph L. (1972). “The Estimation of the Lorenz Curve and Gini Index”. In: *The Review of Economics and Statistics* 54.3, pp. 306–316.
- Gordon, David M. (2018). *Reading the Archives as Sources*. DOI: 10.1093/acrefore/9780190277734.013.227.
- Haas, Michiel de (2022). “Reconstructing income inequality in a colonial cash crop economy: five social tables for Uganda, 1925–1965”. In: *European Review of Economic History* 26.2, pp. 255–283.
- Haas, Michiel A de (2017). “Rural livelihoods and agricultural commercialization in colonial Uganda: Conjunctures of external influences and local realities”. In:
- Hundenborn, Janina, Ingrid Woolard, and Jon Jellema (2019). “The effect of top incomes on inequality in South Africa”. In: *International Tax and Public Finance* 26.
- Jenkins, Stephen P (2017). “Pareto Models, Top Incomes and Recent Trends in UK Income Inequality”. In: *Economica* 84, pp. 261–289.
- Kakwani, Nanak Chand (1980). *Income inequality and poverty : methods of estimation and policy applications*. Tech. rep. World Bank Research Publication 10092.
- Kopczuk, Wojciech (2015). “What Do We Know about the Evolution of Top Wealth Shares in the United States?” In: *Journal of Economic Perspectives* 29.1, pp. 47–66.
- Kuznets, Simon (1955). “Economic growth and income inequality”. In: *The American Economic Review* 45.1, pp. 1–28.
- Lindert, Peter H and Jeffrey G Williamson (1982). “Revising England’s social tables 1688–1812”. In: *Explorations in economic history* 19.4, pp. 385–408.
- Milanovic, Branko (2006). “An estimate of average income and inequality in Byzantium around year 1000”. In: *Review of Income and Wealth* 52.3, pp. 449–470.

- Milanovic, Branko (2009). “Global inequality and the global inequality extraction ratio: the story of the past two centuries”. In: *World Bank Policy Research Working Paper* 5044.
- (2018). “Towards an explanation of inequality in premodern societies: the role of colonies, urbanization, and high population density”. In: *The Economic History Review* 71.4, pp. 1029–1047.
- (2022). “The three eras of global inequality, 1820-2020 with the focus on the past thirty years”. In:
- Milanovic, Branko, Peter H Lindert, and Jeffrey G Williamson (2007). *Measuring Ancient Inequality*. Tech. rep. World Bank Policy Research Working Paper 4412.
- (2011). “Pre-industrial inequality”. In: *The Economic Journal* 121.551, pp. 255–272.
- Modalsli, Jørgen (2015). “Inequality in the very long run: inferring inequality from data on social groups”. In: *The Journal of Economic Inequality* 13.2, pp. 225–247.
- Morrisson, Christian and Wayne Snyder (2000). “The income inequality of France in historical perspective”. In: *European Review of Economic History* 4.1, pp. 59–83.
- Pareto, Vilfredo (1897). *Cours d’Economie Politique, Professé a l’Université de Lausanne. Tome Second*. Lausanne: F. Rouge.
- Piketty, Thomas (2018). *Capital in the twenty-first century*. Harvard University Press.
- Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2022). “Twenty Years and Counting: Thoughts about Measuring the Upper Tail”. In: *The Journal of Economic Inequality* 20.1, pp. 255–264. DOI: 10.1007/s10888-022-09536-8.
- Saez, Emmanuel and Gabriel Zucman (2016). “Wealth Inequality in the United States since 1913: Evidence From Capitalized Income Tax Data”. In: *Quarterly Journal of Economics* 131.2.
- Saito, Osamu (2015). “Growth and inequality in the great and little divergence debate: a Japanese perspective”. In: *The Economic History Review* 68.2, pp. 399–419.
- Shorrocks, Anthony and Guanghua Wan (2008). *Ungrouping income distributions: Synthesising samples for inequality and poverty analysis*. Tech. rep. WIDER Research Paper.
- Sokoloff, Kenneth L and Stanley L Engerman (2000). “Institutions, factor endowments, and paths of development in the new world”. In: *Journal of Economic perspectives* 14.3, pp. 217–232.
- Van Bavel, Bas JP and Jan Luiten Van Zanden (2004). “The jump-start of the Holland economy during the late-medieval crisis, c. 1350–c. 1500 1”. In: *The Economic History Review* 57.3, pp. 503–532.
- Van Ourti, Tom and Philip Clarke (2011). “A simple correction to remove the bias of the Gini coefficient due to grouping”. In: *Review of Economics and Statistics* 93.3, pp. 982–994.
- Van Zanden, Jan Luiten and Bas Van Leeuwen (2012). “Persistent but not consistent: The growth of national income in Holland 1347–1807”. In: *Explorations in economic history* 49.2, pp. 119–130.
- von Fintel, Dieter and Anna Orthofer (2020). “Wealth inequality and financial inclusion: Evidence from South African tax and survey records”. In: *Economic Modelling* 91, pp. 568–578. DOI: <https://doi.org/10.1016/j.econmod.2020.02.001>.

Williamson, Jeffrey G (2010). “Five centuries of Latin American income inequality”. In: *Revista de Historia Economica-Journal of Iberian and Latin American Economic History* 28.2, pp. 227–252.

A Derivation of changes in social table estimates when classes are sub-divided

Alvaredo (2011)'s "two-group" formulation in equation 1 focuses primarily on the contribution of a small top class to the Gini coefficient. However, his analysis also has implications for other parts of the distribution and can be generalized to examine the role of large bottom classes in determining inequality. His simple formulation suggests the possibility that even small changes in G^* – which could arise from creating large bottom classes – can have a significant influence on total inequality if the income share of the bottom of the distribution $B = (1 - S)$ is large. Together $G^*(1 - S)$ can help determine the total *Gini* if $B = (1 - S)$ grows large.

We generalize this approach, but focus on the bottom class. We also incorporate $k > 2$ groups. Our starting point is Dagum's (1997) decomposition of the Gini coefficient into between- and within-group inequality. Firstly the Gini coefficient between group j and h is

$$G_{jh} = \frac{\sum_{i=1}^{N_j} \sum_{r=1}^{N_h} |y_{ij} - y_{hr}|}{N_j N_h (\mu_j + \mu_h)} \quad (\text{A.1})$$

while the Gini coefficient within class j is

$$G_{jj} = \frac{\sum_{i=1}^{N_j} \sum_{r=1}^{N_j} |y_{ji} - y_{jr}|}{2N_j^2 \mu_j} \quad (\text{A.2})$$

The population share of class j is

$$P_j = \frac{N_j}{N} \quad (\text{A.3})$$

and the income share of class j is

$$S_j = \frac{N_j \mu_j}{N \mu} \quad (\text{A.4})$$

$$\Rightarrow \frac{S_j}{P_j} = \frac{\frac{N_j \mu_j}{N \mu}}{\frac{N_j}{N}} = \frac{\mu_j}{\mu} \quad (\text{A.5})$$

$$\Rightarrow S_j P_j = \frac{N_j^2 \mu_j}{N^2 \mu} = P_j^2 \frac{\mu_j}{\mu} \quad (\text{A.6})$$

$$\Rightarrow S_j P_h = \frac{N_j N_h \mu_j}{N^2 \mu} = P_j P_h \frac{\mu_j}{\mu} \quad (\text{A.7})$$

The Gini coefficient for the whole population can be decomposed into within- and between-group components, as per Dagum (1997):

$$G = \underbrace{\sum_{j=1}^k G_{jj} P_j S_j}_{\text{within}} + \underbrace{\sum_{j=1}^k \sum_{h=1}^{j-1} G_{jh} (P_j S_h + P_h S_j)}_{\text{between}} \quad (\text{A.8})$$

Using equations A.4 to A.7, a typical between-group term can be simplified and used to derive G in terms

of class means and population shares – indicators that are the building blocks of social tables:

$$G_{jh}(P_j S_h + P_h S_j) = \frac{\sum_{i=1}^{N_j} \sum_{r=1}^{N_h} |y_{ij} - y_{hr}|}{N_j N_h (\mu_j + \mu_h)} (P_j S_h + P_h S_j) \quad (\text{A.9})$$

$$= \frac{\mu_h - \mu_j}{\mu_j + \mu_h} (P_j S_h + P_h S_j) \text{ because non-overlapping classes entails that } y_{hr} > y_{ij} \text{ for } h > j \quad (\text{A.10})$$

$$= \frac{\mu_h - \mu_j}{\mu_j + \mu_h} P_j P_h \left(\frac{S_h}{P_h} + \frac{S_j}{P_j} \right) \quad (\text{A.11})$$

$$= \frac{\mu_h - \mu_j}{\mu_j + \mu_h} P_j P_h \left(\frac{\mu_h + \mu_j}{\mu} \right) \quad (\text{A.12})$$

$$= P_j P_h \frac{\mu_h - \mu_j}{\mu} \quad (\text{A.13})$$

$$\Rightarrow G = \sum_{j=1}^k G_{jj} P_j S_j + \sum_{j=1}^k \sum_{h=1}^{j-1} P_j P_h \left(\frac{\mu_h - \mu_j}{\mu} \right) \quad (\text{A.14})$$

We now use the decomposition of the Gini coefficient in equation A.14 to compare two scenarios. The first, scenario “ L ”, groups the distribution into k classes. The bottom class constitutes a large proportion of the population (high P_1^L). There is enough (unobserved) variation in this group to generate a substantial unobserved within-group Gini coefficient (G_{11}^L). We use the subscript L for this scenario to correspond to a large bottom class. It follows from equation A.14 that the true Gini coefficient can be expressed as a function of the “group” statistics associated with the k classes:

$$G = \sum_{j=1}^k G_{jj}^L P_j^L S_j^L + \sum_{j=1}^k \sum_{h=1}^{j-1} P_j^L P_h^L \left(\frac{\mu_h^L - \mu_j^L}{\mu^L} \right) \quad (\text{A.15})$$

However, the social tables estimate for scenario “ L ” ignores within-group variation ($G_{jj}^L = 0$):

$$\widehat{G}_L = \sum_{j=1}^k \sum_{h=1}^{j-1} P_j^L P_h^L \left(\frac{\mu_h^L - \mu_j^L}{\mu^L} \right) \text{ where } \mu^L = \sum_{j=1}^k P_j^L \mu_j^L. \quad (\text{A.16})$$

The second scenario extends the analysis to $k + 1$ classes. The large bottom class is sub-divided into two smaller units – hence our use of the superscript “ S ” – with a (slightly) lower population share in the bottom class (P_1^S) and less variation in the bottom class ($G_{11}^S < G_{11}^L$). As before, the “true” Gini can be expressed as a function of the $k + 1$ groups’ attributes:

$$G = \sum_{j=1}^{k+1} G_{jj}^S P_j^S S_j^S + \sum_{j=1}^{k+1} \sum_{h=1}^{j-1} P_j^S P_h^S \left(\frac{\mu_h^S - \mu_j^S}{\mu^S} \right) \quad (\text{A.17})$$

and, with $G_{jj}^S = 0$ the social tables estimate in this scenario reduces to

$$\widehat{G}_S = \sum_{j=1}^{k+1} \sum_{h=1}^{j-1} P_j^S P_h^S \left(\frac{\mu_h^S - \mu_j^S}{\mu^S} \right) \text{ where } \mu^S = \sum_{j=1}^{k+1} P_j^S \mu_j^S. \quad (\text{A.18})$$

Comparing the two biased estimates \widehat{G}_L and \widehat{G}_S reveals how subdividing the bottom class into smaller units influences social tables inequality estimates. The difference between the estimates is determined by

how between-group dynamics change between the two scenarios. We first assume that the estimated mean of the distribution stays constant between the two scenarios. ($\mu^S = \mu^L$). Then:

$$\widehat{G}_S - \widehat{G}_L = \sum_{h=2}^{k+1} P_1^S P_h^S \frac{\mu_h^S - \mu_1^S}{\mu^S} + \sum_{h=3}^{k+1} P_2^S P_h^S \frac{\mu_h^S - \mu_2^S}{\mu^S} - \sum_{h=2}^k P_1^L P_h^L \frac{\mu_h^L - \mu_1^L}{\mu^L} \quad (\text{A.19})$$

$$= \sum_{h=2}^{k+1} P_1^S P_h^S \frac{\mu_h^S - \mu_1^S}{\mu^S} + \sum_{h=3}^{k+1} (P_2^S - P_1^L) P_h^S \frac{\mu_h^S - \mu_2^S}{\mu^S} \text{ because } \mu_{h+1}^S = \mu_h^L \text{ and } P_{h+1}^S = P_h^L \text{ for } h \geq 3 \quad (\text{A.20})$$

$$= \sum_{h=2}^{k+1} P_1^S P_h^S \frac{\mu_h^S - \mu_1^S}{\mu^S} + \sum_{h=3}^{k+1} (P_2^S - P_1^S - P_2^S) P_h^S \frac{\mu_h^S - \mu_2^S}{\mu^S} \quad (\text{A.21})$$

$$= P_1^S P_2^S \frac{\mu_2^S - \mu_1^S}{\mu^S} > 0 \text{ because } \mu_2^S > \mu_1^S \text{ and } \mu_1^S > 0; \mu_2^S > 0; P_1^S > 0; P_2^S > 0 \quad (\text{A.22})$$

The two scenarios give different inequality estimates. Specifically, equation A.22 shows that scenario “S” yields a larger Gini coefficient than scenario “L”, as expected. A large bottom class underestimates the Gini coefficient. The difference comes from the fact that there is significant variation in the scenario “L” bottom class. When it is sub-divided into two groups, mean incomes of the sub-divided bottom classes differ enough to change the overall between-group social table estimate. The Gini coefficient can even be sensitive to small mean differences if the population share of one or both of the groups remains “large”. In contrast, if either $P_1^S \rightarrow 0$ or $P_2^S \rightarrow 0$, $\widehat{G}_S - \widehat{G}_L \rightarrow 0$, as expected.

The problem of large bottom classes can also be analysed from the within group perspective. If $\widehat{G}_S - \widehat{G}_L > 0$, sub-division recovers a part of the unobserved within-group variation by estimating \widehat{G}_S . Changes in within-group variation are not observed, but mirror changes in observed between-group variation. Analysing these theoretical changes gives additional stylised facts about sub-division. Noting that the “top” groups have not changed across scenarios “S” and “L”:

$$\widehat{G}_L - \widehat{G}_S = G_{11}^S P_1^S S_1^S + G_{22}^S P_2^S S_2^S - G_{11}^L P_1^L S_1^L \quad (\text{A.23})$$

$$= G_{11}^S P_1^S S_1^S + G_{22}^S P_2^S S_2^S - G_{11}^L (P_1^S + P_2^S) (S_1^S + S_2^S) \quad (\text{A.24})$$

$$= G_{11}^S P_1^S S_1^S + G_{22}^S P_2^S S_2^S - G_{11}^L P_1^S S_1^S - G_{11}^L P_1^S S_2^S - G_{11}^L P_2^S S_1^S - G_{11}^L P_2^S S_2^S \quad (\text{A.25})$$

$$= (G_{11}^S - G_{11}^L) P_1^S S_1^S + (G_{22}^S - G_{11}^L) P_2^S S_2^S - G_{11}^L (P_1^S S_2^S + P_2^S S_1^S) \quad (\text{A.26})$$

Because variation within either of the sub-divided groups (G_{jj}^S $j = 1, 2$) is smaller than the variation in the “large” bottom class (G_{11}^L), ($G_{ii}^S - G_{11}^L$) < 0 and the social tables inequality estimate (for the whole distribution) increases ($\widehat{G}_L - \widehat{G}_S$) < 0 . Even small within-group variations in the bottom tail can produce meaningful differences in inequality estimates if the bottom class is sub-divided. This is because the change is pronounced if either of the sub-divided classes has a significantly large population (P_j^S) or income share (S_j^S). Even incremental sub-divisions of the bottom class could produce significantly different Gini coefficients.

B Replication results

Table B.1: Replication of Milanovic (2018) using various samples

<i>Sample</i>	(1)	(2)	(3)	(4)	(5)	(6)
	Milanovic (2018)			All		
Population Density (<i>people/km²</i>)	-0.001* (0.000)	-0.001* (0.000)	-0.001* (0.000)	-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)
GDP (<i>per capita 1990 PPP</i>)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)
GDP ²	0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Number of Classes			-0.000 (0.002)			0.001 (0.001)
Constant	0.437*** (0.104)	0.442*** (0.113)	0.441*** (0.116)	0.478*** (0.026)	0.485*** (0.027)	0.485*** (0.028)
Incl micro estimates	Y	N	N	Y	N	N
<i>N</i>	33	26	26	80	73	67

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

C Sources used to compile social tables data

Table C.2: Sources used to compile dataset

Social Tables List			
Paper	Publication year	Authors	Countries
Long-run trends in economic inequality: lessons from colonial Botswana (1921 to 1974)	2016	Bolt and Hillbom	Botswana
Preindustrial Inequality	2011	Milanovic, Lindert & Williamson	Various countries
Between the colonial heritage and the first globalisation boom: On income inequality in the southern cone	2010	Willebald	Various countries
American Incomes before and After the Revolution	2013	Lindert & Williamson	USA
Income Inequality in Colonial Africa: Building Social Tables for Pre-Independence Central African Republic, Ivory Coast, and Senegal	2019	Alfani & Tadei	Senegal, Ivory Coast
Towards an explanation of inequality in premodern societies: the role of colonies, urbanisation, and high population density	2018	Milanovic	Various countries
Economic Inequality in Ghana 1891 to 1960	2021	Bolt & Aboagye	Ghana
Reconstructing income inequality in a colonial cash crop economy: five social tables for Uganda, 1925–1965	2021	de Haas	Uganda
Income Inequality in Mexico 1895 -1940: Industrialization, revolution, Institutions	2019	Garza & Bengtsson	Mexico
Living costs, real incomes and inequality in colonial Jamaica	2019	Burnard, Panza & Williamson	Jamaica
Income Distribution in Rural Buenos Aires, 1839-1867	2015	Bertola, Gellman & Santilli	Argentina
International Trade, Government, and Income Distribution in Peru Since 1870	1990	Berry	Peru
Economic Development and Fluctuations in Earnings Equality in the Very Long Run: The Evidence from Latin America 1900–2000	2008	Fitzgerald	Various countries

Lund Papers in Economic History
ISRN LUSADG-SAEH-P--23/247--SE+36

© The author(s), 2023

Orders of printed single back issues (no. 1-65)

Department of Economic History, Lund University
Postal address: P.O. Box 7083, S-220 07 Lund, Sweden
Telephone: +46 46 2227475
Telefax: +46 46 131585

Full-text electronic issues (no. 58, 60, 61, 66--)

www.ekh.lu.se

Lund Papers in Economic History are published by the Department of Economic History, Lund University, Sweden. This series replaces the former series under the title *Meddelande från ekonomisk-historiska institutionen, Lunds universitet*. The change of name reflects the orientation of the series towards an international readership. The series is multilingual, but the majority of the working papers appear in English.

Lund Papers in Economic History are published as occasion arises, not at fixed intervals. Printed issues are distributed to libraries. From 1999 and onwards, full-text electronic issues are also available on www.ekh.lu.se. Those who would be interested in receiving information by email on new issues of **Lund Papers in Economic History** are requested to send an email message to Lund.Papers@ekh.lu.se.



LUND
UNIVERSITY