



The Determinants of Incomes and Inequality: Evidence from Poor and Rich Countries

Christoph Lakner

A thesis submitted for the degree of Doctor of Philosophy in Economics

St. Antony's College
University of Oxford

Trinity Term 2014

Abstract

This thesis consists of four separate chapters which address different aspects of inequality and income determination. The first three chapters are country-level studies which examine (1) how incomes are shaped by spatial price differences, (2) the factor income composition, and (3) enterprise size. The final chapter analyses how income inequality changed at the global level.

The first chapter investigates the implications of regional price differences for earnings differentials and inequality in Germany. I combine a district-level price index with administrative earnings data from social security records. Prices have a strong equalising effect on district average wages in West Germany, but a weaker effect in East Germany and at the national level. The change in overall inequality as a result of regional price differences is small (although significant in many cases), because inequality is mostly explained by differences within rather than between districts.

The second chapter is motivated by the rapid increase in top income shares in the United States since the 1980s. Using data derived from tax filings, I show that this pattern is very similar after controlling for changes in tax unit size. Over the same period as top income shares increased, the composition of these incomes changed dramatically, with the labour share rising. Using a non-parametric copula framework, I show that incomes from labour and capital have become more closely associated at the top. This association is asymmetric such that top wage earners are more likely to also receive high capital incomes, compared with top capital income recipients receiving high wages.

In the third chapter, I investigate the positive cross-sectional relationship between enterprise size and earnings using panel data from Ghana. I find evidence for a significant firm size effect in matched firm-worker data and a labour force panel, even after controlling for individual fixed effects. The size effect in self-employment is stronger in the cross-section, but it is driven by individual time-invariant characteristics.

The final chapter studies the global interpersonal income distribution using a newly constructed and improved database of national household surveys between 1988 and 2008. The chapter finds that the global Gini remains high and approximately unchanged at around 0.7. However, this hides a substantial change in the global distribution from a twin-peaked distribution in 1988 into a single-peaked one now. Furthermore, the regional composition of the global distribution changed, as China graduated from the bottom ranks. As a result of the growth in Asia, the poorest quantiles of the global distribution are now largely from Sub-Saharan Africa. By exploiting the panel dimension of the dataset, the analysis shows which decile-groups within countries have benefitted most over this 20-year period. In addition, the chapter presents a preliminary assessment of how estimates of global inequality are affected by the likely underreporting of top incomes in surveys.

Word Count and Format

The computerised word count is 78,017. Taking page 2 as a representative page, the thesis counts 376 words per page. Multiplying this by 207 pages (excluding bibliography) results in an overall word count of 77,832.

This thesis consists of four main chapters, an introduction and a conclusion. The chapters contain a number of sections and an appendix. All relevant tables and figures are at the end of each chapter.

Declaration

Chapter 3 is a revised version of the first part of my MPhil thesis, which was submitted in Trinity Term 2010 as part of the requirements for the degree of MPhil in Economics, University of Oxford. The title of my MPhil thesis was “Enterprise size and selection as determinants of urban income: evidence from Ghana”.

Chapter 4 is co-authored with Branko Milanovic, formerly at the World Bank. A shorter and preliminary version of this chapter appeared as World Bank Policy Research Working Paper 6719.

Section 2.5 of Chapter 2 is my contribution to a joint working paper with Rolf Aaberge, Sir Tony Atkinson and Sebastian Königs (Aaberge et al., 2013).

Acknowledgements

This thesis would not have been possible without the support and encouragement from a large number of people. First and foremost, I would like to thank my supervisors Sir Tony Atkinson and Francis Teal for their invaluable help and guidance throughout this research. I am very grateful for their patience and I have learnt tremendously from both of them. I thank Emmanuel Saez for supervising me when I was a visiting student at the University of California, Berkeley, and for his assistance with the US tax data. I am very grateful to Branko Milanovic for being such an enthusiastic co-author and sharing his expertise on global inequality. I would also like to thank my undergraduate tutor John Quah for his inspiration to pursue graduate studies. I thank my examiners Adrian Wood and Stephan Klasen for many helpful comments.

The Department of Economics at the University of Oxford created a very stimulating environment to undertake this research through the coursework, seminars and workshops. I greatly enjoyed being part of the Centre for the Study of African Economies (CSAE) where I learnt a lot about empirical development economics. I gratefully acknowledge financial support from the Economic and Social Research Council, the Studienstiftung des deutschen Volkes, the CSAE, the Institute for New Economic Thinking (INET) at the Oxford Martin School, the George Webb Medley Fund, the Center for Equitable Growth at UC Berkeley, the German Institute for Employment Research (IAB) and the International Association for Research in Income and Wealth (IARIW).

I would like to thank staff at the IAB, the Federal Institute for Research on Building, Urban Affairs and Spatial Development (BBSR) within the Federal Ministry for Environment and Building, and the German Federal Statistics Office for providing access to the various datasets, answering numerous questions and for their hospitality. I am grateful to Johannes Ludsteck for sharing the imputation code used in Dustmann et al. (2009). I thank Francis Teal and Paolo Falco at the CSAE for their help with the Ghana data. I also thank Statistics Finland, Statistics Portugal and Eurostat for providing income statistics.

I would like to extend my thanks to seminar participants of the Gorman workshop at the University of Oxford, the informal INET workshop, the World Bank, the Studienstiftung's Doktorandenforum, the Singapore Economic Review Conference 2011, the CSAE Conferences 2012 and 2014, the IARIW General Conference 2012, the Kiel IfW Centenary Conference 2014 and the Royal Economic Society Conference 2014. I am grateful for helpful comments received from Andrea Brandolini, Shaohua Chen, David Card, Paolo Falco, Chico Ferreira, Bronwyn Hall, Philippe van Kerm, Patrick Kline, Sebastian Königs, Peter Lanjouw, Salvatore Morelli, Enrico Moretti, John Muellbauer, John Pencavel, Simon Quinn, Justin Sandefur, Marcel Timmer and Till von Wachter.

I thank my friends inside and outside the Department of Economics for their support. I am grateful to Rebecca Scott for her invaluable help with the submission. In particular, I wish to thank La-Bhus Fah Jirasavetakul for all her support and wisdom over the years. Finally, I wish to thank my parents and sisters for their support and encouragement.

Contents

Introduction	1
1 The Inequality of Real Wages in Germany	4
1.1 Literature review	7
1.2 The price data	11
1.3 The wage data	15
1.4 Constructing real earnings and summary statistics	19
1.5 Results at the district level	20
1.6 Results at the individual level	23
1.7 Conclusion	25
1.8 Main Tables and Figures	26
1.A Appendix Tables and Figures	35
2 Top Incomes in the USA, 1960-2005: Family Size and Factor Income Composition	36
2.1 Literature review and motivation	39
2.2 The data and methodology	41
2.2.1 Unit of analysis	42
2.2.2 Income concept	44
2.2.3 Income shares as a measure of inequality and control totals . . .	46
2.2.4 Summary	48
2.2.5 Summary statistics	49
Part I Top income shares adjusted for tax unit size	51
2.3 Adjusting for tax unit size	51
2.3.1 Empirical specification	52
2.3.2 Results	55
Part II Association between labour and capital income	58
2.4 Factor income decomposition: The alignment coefficient	60
2.4.1 Results	63
2.5 The copula function	67
2.5.1 Association matrices as a discrete approximation	69
2.5.2 Differences between cumulative association matrices	72
2.6 Conclusion	75
2.7 Main Tables and Figures	77
2.A Appendix	88
2.A.1 Tax evasion	88
2.A.2 Changes in the tax code over time	89
2.A.3 Effects of changes in the tax code	91
2.A.4 Comparing market income with Haig-Simons income	95
2.A.5 Using before-tax incomes	97

2.B	Appendix Tables	99
3	Enterprise Size and Earnings: Evidence from Urban Ghana	102
3.1	Econometric framework	104
3.1.1	Literature review	104
3.1.2	Estimation	108
3.2	Data description	110
3.2.1	Summary statistics	113
3.3	Results	114
3.3.1	The effect of scale in wage-employment	114
3.3.2	The size effect in self-employment	117
3.4	Conclusion	118
3.5	Main Tables and Figures	120
3.A	Appendix Tables and Figures	125
4	Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession	129
4.1	Literature review and our approach	134
4.1.1	Global inequality	135
4.1.2	Global poverty	138
4.1.3	PPP exchange rates	139
4.1.4	National accounts and household surveys	145
4.1.5	Top incomes	146
4.1.6	Missing top incomes and the national accounts discrepancy	148
4.2	Data construction and methodology	150
4.2.1	Data source	150
4.2.2	Survey selection and the definition of regions	151
4.2.3	Welfare concept	153
4.2.4	Pareto imputation and scaling to national accounts consumption	154
4.2.5	The panel dimension of our data	156
4.3	The cross-sectional distribution over time	157
4.3.1	Summary statistics	157
4.3.2	The Gini index of the global distribution of income	158
4.3.3	Alternative measures of inequality	161
4.3.4	Regional inequality and between-country decomposition	162
4.3.5	Growth incidence curves	165
4.3.6	The regional composition of the global distribution	167
4.4	Accounting for top incomes in the computation of global inequality	170
4.5	Changes over time: Who has grown fastest and slowest?	177
4.5.1	Quasi-non-anonymous growth incidence curves	177
4.5.2	Most successful and least successful	180
4.6	Conclusion	184
4.7	Main Tables and Figures	188
4.A	Appendix	199
4.B	Appendix Tables	201
	Conclusion	205
	Bibliography	208

List of Tables

1.1	Regression of the price index on rural/urban and state indicators	28
1.2	Summary statistics for different wage variables in 2007, nominal and real .	28
1.3	The 10 richest and poorest districts and their prices	32
1.4	Relationship between median and mean wages and prices across districts .	32
1.5	Effect of prices on inequality in 2007 by region.	34
1.A.1	Number and proportion of censored observations in 2007	35
2.1	Summary statistics	78
2.2	Association matrix: Perfect association (frequencies in %).	82
2.3	Association matrix: Independence (frequencies in %).	82
2.4	Association matrix: 1980 (frequencies in %)	83
2.5	Association matrix: 1990 (frequencies in %)	83
2.6	Association matrix: 2000 (frequencies in %)	83
2.7	Inverse cumulative association matrix: 1980 (frequencies in %).	84
2.8	Inverse cumulative association matrix: 1990 (frequencies in %).	84
2.9	Inverse cumulative association matrix: 2000 (frequencies in %).	84
2.10	Conditional distribution for Top 1% (frequencies in %).	86
2.11	Change in inverse cumulative association matrix: 2000 and 1980 (in pp). .	86
2.12	Change in inverse cumulative association matrix: 1990 and 1980 (in pp). .	87
2.13	Change in inverse cumulative association matrix: 2000 and 1990 (in pp). .	87
2.B.1	Total income share excl. capital gains (in %): Comparing to Saez (2013) .	99
2.B.2	Total income share incl. capital gains (in %): Comparing to Saez (2013). .	100
2.B.3	Total income share incl. capital gains (in %): Comparing Part II sample .	101
3.1	Matched firm-worker data: Summary statistics by firm size	120
3.2	Labour force data: Summary statistics by occupation and size	121
3.3	Matched firm-worker data: Baseline results	122
3.4	Matched firm-worker data: Difference- and System-GMM	122
3.5	Labour force data, wage-employees: Baseline results	123
3.6	Labour force data, wage-employees: Difference- and System-GMM	123
3.7	Labour force data, self-employed: Baseline results.	124
3.8	Labour force data, self-employed: Difference- and System-GMM.	124
3.A.1	Matched firm-worker data: Distribution of firm size by year	126
3.A.2	Matched firm-worker data: Change in firm size since last year	126
3.A.3	Labour force data: Distribution of firm size by year.	127
3.A.4	Labour force data: Change in firm size since last year	127
3.A.5	Labour force data: Distribution of number of employees by year.	127
3.A.6	Labour force data: Change in no. employees since last year.	127
3.A.7	Matched firm-worker data: Analysis for men only	128

3.A.8	Matched firm-worker data: Analysis for women only	128
3.A.9	Labour force data, wage-employees: by gender	128
3.A.10	Labour force data, self-employed: by gender	128
4.1	Sample summary statistics	188
4.2	Panel summary statistics: Number of countries by panel duration	189
4.3	Global and regional inequality	190
4.4	Robustness check on global Gini index: Accounting for top incomes.	191
4.5	20 biggest gainers and losers: 1988-2008	196
4.B.1	Comparison with previous estimates of global inequality.	201
4.B.2	Global Gini index (%) for balanced sample	202
4.B.3	Size of national accounts adjustment: Survey mean to max(survey mean, national accounts consumption) (%)	202
4.B.4	Comparison of Pareto constant with fiscal data	203
4.B.5	Comparing top-heavy shares with the World Top Incomes Database	203
4.B.6	20 biggest gainers and losers: 1993-2008	204

List of Figures

1.1	Regional price index (BBSR, 2009)	26
1.2	Density of price index	27
1.3	Regional composition of the distribution of median wages	27
1.4	Regional distribution of nominal and real median wages per day (by decile); wages imputed above 169, 146	29
1.5	Median wages and prices across districts: National, West and East	30
1.6	Mean wages and prices across districts: National, West and East	31
1.7	Comparing distribution of prices, median and mean wages by region	33
1.A.1	Nominal earnings for Berlin, East and West Germany: Raw wages	35
2.1	Income share of the Top 1% of income recipients	77
2.2	Income composition of the Top 1% of income recipients	77
2.3	Accounting for tax unit size (Top 0.5% and Top 1%)	79
2.4	Accounting for tax unit size (Top 5% and Top 10%)	80
2.5	Factor income decomposition of Top 1% share	81
2.6	Factor income decomposition of Top 5% share	81
2.7	Factor income decomposition of Top 0.5% share	82
2.8	Depiction of the inverse cumulative association matrix: Top 1%	85
2.9	Depiction of the inverse cumulative association matrix: Top 5%	85
2.10	Depiction of the inverse cumulative association matrix: Top 0.5%	86
3.1	Matched firm-worker data: Distribution of earnings by size	120
3.2	Labour force data: Distribution of earnings by size and sector	121
3.A.1	GDP and firm size across countries	125
3.A.2	GDP and wage employment across countries	125
3.A.3	Comparing distribution of firm size in matched and labour force data	126
4.1	Comparison of global and within-country inequality	189
4.2	Lorenz curves in 1988 and 2008	191
4.3	Global growth incidence curve (for fractile groups)	192
4.4	The global distribution of income	193
4.6	The regional composition of the global distribution of income	193
4.5	The changing global distribution of income	194
4.7	Comparing different Gini coefficients	195
4.8	Baseline and top-heavy distribution of income	195
4.9	Quasi-non-anonymous growth incidence curve	197
4.10	Positional gains	198

Introduction

This thesis addresses four different aspects of inequality and the determination of income. In this introduction, I briefly outline the motivation for each chapter. While each of the chapters is a standalone piece, I also attempt to recognise the similarities and clarify the differences.

In the first chapter, I examine whether failing to account for spatial price differences distorts earnings differentials and inequality across Germany. I use a new district-level price index in combination with administrative earnings data which have the necessary sample size for a disaggregated regional analysis. More than 20 years after the fall of the Berlin Wall, there is a strong public interest in whether real wages have converged between East and West Germany. Furthermore, I am interested in disparities within East and West Germany given its federal structure and by being a large and diverse European country.

The second chapter is motivated by the rapid increase in top income shares in the USA that has been documented by Piketty and Saez (2003) and others. I use data from tax records which provide more accurate information about the very top tail of the distribution. The existing literature measures the inequality amongst tax units without an adjustment for tax unit size, which differs from the approach usually adopted in a household survey based analysis. In the first part of the chapter I control for tax unit size using the number of exemptions claimed on the tax return. The second part of the chapter examines the joint distribution of capital and labour incomes. The previous literature showed that the income share from labour increased at the top. Using a non-parametric copula framework, I investigate whether this is driven by changes in

the marginal distributions of capital and labour, or by the association between the two income sources.

In the third chapter, I attempt to explain the observed variation in labour market earnings in urban Ghana. In particular, I examine why enterprise size has such a large positive effect on earnings in both wage- and self-employment. I estimate the size effect separately in the two types of employment, as I expect the mechanisms to be quite different. I use a matched firm-worker panel dataset, which oversamples large firms, and a representative labour force panel survey. In these two panel datasets I examine the roles of observed and unobserved time-invariant characteristics, and use dynamic panel estimators as additional controls for endogeneity.

The final chapter takes the concern for inequality from the national to the global level. I am interested in measuring inequality amongst all individuals in the world regardless of their country of residence. The chapter uses a new database of national household surveys between 1988 and 2008. I present new estimates of global inequality and a first adjustment for the likely underreporting of top incomes in household surveys, as found by the top income literature reviewed in Chapter 2. I analyse the changes in the regional composition that underlie the trends in overall inequality. Furthermore, I exploit the structure of the data as a 20-year panel of country-decile-groups.

The four chapters differ in a number of ways. First, the sources of data are quite different and have been chosen with the particular application in mind. Both the Germany and US chapters (Chapters 1 and 2) use administrative data. The German data are based on social security records and miss the top of the distribution because they are top censored and exclude the self-employed. On the other hand, the US data are from tax filings and thus use taxable income which is likely to understate incomes at the bottom. Furthermore, the US data deliberately oversample the top of the distribution which make them especially useful for studying top incomes. The Ghana chapter (Chapter 3) uses two panel datasets which allow me to control for time-invariant characteristics. These datasets are both relatively small surveys which were collected by the Centre for the Study of African Economies at the University

of Oxford. The global chapter (Chapter 4) combines national household surveys from around the world which attempt to be representative of the entire population. However, as discussed in that chapter, it is quite likely that they understate top incomes.

Second, these data sources employ different income measures. The German chapter (chapter 1) uses labour earnings. This is quite similar to the Ghana chapter (chapter 3), except that it includes the self-employed which were excluded from the German data. The US chapter (Chapter 2) combines (taxable) earned income and capital income and analyses the joint distribution of the two income sources. I explain this income definition in detail in that chapter, as it partly depends on the tax code. In the global chapter (Chapter 4), income is defined as disposable income.

Third, the unit of analysis in the chapters on Germany and Ghana (Chapters 1 and 3) is an individual. In the global chapter (Chapter 4), disposable income is measured at the household-level but it is expressed in household per capita terms, so the unit of analysis is still the individual. Following the related literature, the US chapter (Chapter 2) mainly measures the inequality amongst tax units, but in the first part of that chapter I analyse the effect of controlling for tax unit size.

Fourth, spatial price differences are taken into account in the Germany and global chapters (Chapters 1 and 4) but in a different way. In Chapter 1, regional price differences within Germany are the central theme of the analysis. Chapter 4 largely ignores price differences within countries but adjusts for international differences by using purchasing power parity (PPP) exchange rates.¹ While the German price index ignores spatial variation in consumption baskets, this would clearly be untenable for an international comparison with PPP exchange rates.

Each chapter is concerned with a theoretical question and the imperfect empirical evidence available to answer that question. In each case, I discuss in detail just what one can and cannot say. Furthermore, every chapter attempts to deepen our understanding of inequality and to bring out the subtleties of the term - subtleties which are often overlooked when people simply talk about “income inequality”.

¹For the three most populous developing countries, China, India and Indonesia, I account for price differences between urban and rural areas.

Chapter 1

The Inequality of Real Wages in Germany¹

Introduction

This chapter investigates how accounting for regional price variation affects regional earnings disparities and aggregate inequality in Germany. There is substantial inter-regional variation in both nominal earnings and prices, and I am interested in how these two variables interact to produce differences in real earnings. As it has recently unified two very different economies, Germany offers an interesting case study. 20 years after the fall of the Berlin Wall, it remains to be seen how real earnings compare between and within East and West Germany. Accounting for regional price differences has important policy implications in areas such as social security payouts, national wage bargaining, and the proposed introduction of a comprehensive minimum wage in Germany.

Despite this policy interest, price data particularly at a disaggregated level are rarely available. For example, there exist comprehensive price indices at the level of 12 regions in the UK (Wingfield et al., 2005) and 20 regions of Italy (Pittau et al., 2011) which have

¹This study uses the weakly anonymous Sample of Integrated Labour Market Biographies (Years 2004-2008). Data access was provided via on-site use at the Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB) and subsequently remote data access.

been calculated by the respective National Statistics Offices. Moretti (2013) estimates the spatial variation of housing costs for almost 400 regions in the US. Previous studies on Germany mostly relied on price data for 50 cities. I use a new district-level price index for 2007 published by a division of the German Federal Ministry for Environment and Building (BBSR, 2009). This price index offers an unprecedented disaggregation at the level of almost 400 districts and is very comprehensive in that it covers a large part of the national consumer price index. I am not aware of any other research which uses the full regional information contained in this index to study individual earnings inequality in Germany. Most of the literature has focused on the differential between East and West Germany, which ignores the substantial variation within these regions. In order to study real earnings inequality I combine the price data with earnings information from the Sample of Integrated Labour Market Biographies (SIAB) published by the Institute for Employment Research (IAB) (SIAB, 2010). This dataset has the necessary regional representativeness and sample size.

I find that prices have an equalising effect on district wage differentials. This effect is stronger in West than East Germany because the price distribution in the East is more concentrated than in the West relative to the dispersion of average wages. Across individuals, the change in wage inequality as a result of controlling for regional price differences is small, although significant. The between-district component of this inequality falls by 40% nationally. Consistent with the results for district averages, this effect is weaker in East (around 35%) but stronger in West Germany (almost 60%). The effect on wage inequality is small because wage inequality is mostly explained by differences within rather than between districts.

The motivation for studying real wages is obviously that they reflect the material standard of living, which a given amount of money could buy in different locations. There are a few conceptual issues which I want to highlight at the outset. First, the price index holds the basket of goods constant across localities in order to estimate a pure price effect. This basket corresponds to the one used for estimating the national CPI, i.e. it is an average across all the districts. Therefore, I need to assume that

variation in tastes and needs (e.g. climatic conditions), as well as substitution effects in response to differences in relative prices, are small. Furthermore, I assume that consumption patterns do not vary with income within a district, thus assuming that preferences are homothetic.² Second, I need to argue that differences in prices do not simply compensate for differences in local amenities, such as natural beauty, cultural institutions and access to public services (Aaberge et al., 2010).³ Deaton and Dupriez (2011b) argue that this view might explain the limited attention paid to price differences within countries. This is a difficult conceptual issue which I discuss in more detail in Section 1.1. In brief, I would argue that location choices are not exclusively driven by amenities but also depend on local labour demand, i.e. the location of employment opportunities.⁴

This chapter is structured in seven main sections. First, I review the literature which motivated my analysis and address the role of amenities in more detail. Second, I describe the price data and analyse factors explaining price dispersion. Third, I discuss the earnings data and how I address the right-censoring of the wage variable which closely follows Dustmann et al. (2009). Fourth, I explain the construction of real wages and present summary statistics. Fifth, I analyse the effect of regional price differences at the district-level. Sixth, I examine how district price differentials affect the inequality amongst individuals. The last section concludes.

²Under homothetic preferences, consumers with different incomes have the same expenditure shares when they face the same prices. Muellbauer (1974) shows that consumer expenditure shares typically depend on income. In Germany, the national weights used by the index are dominated by households in the middle and higher income range, but allowing expenditure shares to vary with income produces very similar inflation rates (Breuer and Mehrhoff, 2009). In the USA, Handbury (2013) finds that an index with homothetic preferences describes the consumption expenditures for low- and medium-income households well but breaks down for higher incomes (between 13% and 21% of her sample). In my analysis I am likely to miss high incomes due to censoring in the earnings variable and having a sample which misses the self-employed and civil servants.

³It could also be argued that higher prices in cities compensate for better access to cultural institutions and other public services, such as education or public transport. As a result, the price level in urban areas would be overstated relative to rural areas, since city-dwellers receive additional benefits.

⁴Using the same price data, Deckers et al. (2013) study the effect of regional price levels on life satisfaction. Amenities are of course likely to play an important role for life satisfaction. They do not discuss amenities but they include district fixed effects, which would control for time-invariant amenities. They find that higher prices reduce life satisfaction at the bottom of the distribution. This suggests that amenities do not fully compensate for higher price levels. See also footnote 6.

1.1 Literature review

I briefly review the literature that has studied the implications of spatial price differences for inequality and poverty. I concentrate on spatial price variation and do not address how prices may vary with income (and thus within districts), which has also been the subject of a substantial literature.⁵ The paper most closely related to my analysis is Goebel et al. (2009) who use the same regional price level data in combination with the German Socio-Economic Panel (SOEP).⁶ Given the sample size of the SOEP data, they cannot exploit the full disaggregated price index and restrict their analysis to comparing East and West Germany.⁷ They find that after adjusting for price differences, average real disposable incomes remain approximately one fifth lower in East than West.⁸ Inequality, as measured by the Gini index, falls by less than 0.5pp after accounting for prices differences.⁹ However, they do find significant changes in the regional distribution of poverty; it declines in the East and increases in the West.¹⁰ Using more aggregated price data and the SOEP, Hillringhaus and Peichl (2010) find that accounting for prices increases poverty in Southern Germany.

⁵Also see footnote 2. In India, Rao (2000) documents that the poor face higher unit prices than the rich which is largely explained by differences in the quantities purchased. As a result, accounting for these price differences actually increases inequality. By contrast, Ruiz-Castillo (1998) finds that in Spain, the prices paid by the rich increased faster than those of the poor, and as a result real inequality fell by more than nominal inequality. Using unit values from household surveys, Deaton and Dupriez (2011b) find that accounting for quality differentials reduces the urban-rural price differential because urban households purchase higher-quality goods. Moretti (2013) finds that accounting for prices reduces the increase in the college premium, i.e. the prices of (rich) college graduates increased faster than those of (poor) high school graduates. But the source of variation are local differences which affect high school and college graduates differentially due to location choices. The same applies to my analysis.

⁶As far as I am aware, there is only one other paper using this dataset. Deckers et al. (2013) was released as a working paper substantially after I had completed the first draft of this chapter.

⁷The sample size of the SOEP is too small for a disaggregated analysis. Only seven districts have more than 100 household observations.

⁸By contrast, Roos (2006) concluded that real disposable income had converged between East and West Germany by the mid-2000s. His analysis is based on price level data for 50 cities (Ströhl, 1994) and an econometric extrapolation to the state-level. The price data used by Goebel et al. (2009) and in this chapter are much more reliable. Brück and Peters (2009) also conclude that incomes between East and West Germany converged after controlling for price differences but it is unclear how their price index is constructed. In a very similar paper, they find that average incomes are approximately a fifth lower in the East than the West after adjusting for price differences (Brück and Peters, 2010).

⁹In Italy, the Gini index of disposable income falls by 4.5% after accounting for price differences at the level of 20 regions (Pittau et al., 2011).

¹⁰Similarly, Slesnick (2002) finds only a small effect on inequality, but a more substantial change in poverty, after adjusting for price differences across five large US regions.

With the SIAB (2010) data I cannot analyse poverty because these data do not contain information about non-labour incomes and it is not possible to create households. However, I can extend their analysis in two major ways. First, by using a much larger dataset, I will be able to exploit the full price information for the first time and examine earnings differentials between districts. A comparison between East and West Germany hides a significant amount of variation which exists within these regions. In fact, while the average price level is lower in the East than the West, the least expensive district is not in Eastern Germany but in Northeast Bavaria. Second, I will use decomposable inequality measures to analyse the effect of spatial price differences on interpersonal inequality in more detail and thus examine why the effect on overall inequality is small.

This research is also closely related to Moretti (2013), who finds that accounting for regional price differences reduces the increase in the college wage premium in the US, which could be seen as a measure of inequality.¹¹ A large body of literature has shown an increase in the US college wage premium during the last three decades (Goldin and Katz, 2007), as reviewed in more detail in the chapter on US top incomes. Moretti (2013) shows that the prices faced by college graduates increased faster because their jobs are located in metropolitan areas with high housing costs (mostly in the bicoastal states).¹² In the UK, spatial price differences at the level of 12 regions reduce the graduate-high school wage gap by almost 7%, and shrink the increase between 1997 and 2008 by more than half (Rienzo, 2010).

A number of papers have studied the effect of price differentials on rural-urban poverty rates. Given that prices tend to be higher in urban than rural areas, such an adjustment tends to decrease rural poverty and increase urban poverty. Mogstad et al. (2007) document this pattern in Norway using municipal-level variation in housing costs. In the USA, Jolliffe (2006) finds that after accounting for differences in rental

¹¹I cannot apply his analysis because I have only cross-sectional instead of panel data on prices. Furthermore, the German education system is quite different with its important vocational training sector. For example, Pischke and von Wachter (2008) find a zero return to compulsory schooling in Germany, which contrasts with the findings for other countries.

¹²In a similar analysis, Curran et al. (2008) investigate how poverty rates change across US cities after adjusting for price differences.

costs, the urban poverty rate increases above the rural poverty rate. In Italy, poverty remains largely concentrated in the South even after controlling for regional price differences (Pittau et al., 2011). In the chapter on global inequality, I address this issue in developing countries and for the measurement of global poverty. For example, Ravallion et al. (2007) find that in China and other developing countries urban-rural price differentials are large and can thus have substantial effects on the spatial distribution of poverty.

The role of amenities

A crucial question in interpreting the importance of regional price differences is the role of local amenities. In the original Roback (1982) model, local amenities are completely reflected in land prices, i.e. they are simply compensating differentials for natural beauty or public services.¹³ Workers choose locations depending on their heterogeneous tastes for these amenities. The model assumes free labour mobility which differentiates it from models of international price differentials.¹⁴ According to the Roback view, it would be wrong to adjust down the nominal wage of someone living in a high-price location because the higher price simply reflects the non-monetary benefit of living in a desirable location. Moretti (2013) provides an alternative explanation: The demand for high-skilled (and high-paid) individuals might be greater in some places because of local productivity shocks. In other words, bankers do not particularly like living in London, but their jobs happen to be there. Therefore, the inequality in welfare is

¹³Of course, there could be many other models of price determination. For example, Somekh (2012) explains spatial price differences by the size of the middle class which searches most aggressively for low prices.

¹⁴In general, labour mobility within countries is certainly much higher than across countries but even within countries labour mobility is probably imperfect (e.g. the Hukou system in China (Wang et al., 2011), high home ownership rates in the USA (Blanchflower and Oswald, 2013)). At the international level, the canonical Harrod-Balassa-Samuelson hypothesis assumes that labour is immobile. It is used to explain why the price level (more precisely, the price level of non-tradables) rises with GDP per capita which is sometimes called the “Penn effect” (Samuelson, 1994). In this model, there exist productivity differences across countries in the production of tradable goods, while their price is fixed at the world price. There are no productivity differences in the non-traded sector. In equilibrium, wages in the traded and non-traded sectors have to be equalised within a country (otherwise workers would switch between sectors). The price of non-traded goods (and thus, the overall price level) is higher in countries with higher productivity in tradables. This arises because wages in the non-tradable sector are higher in these countries while the productivity in non-tradables is the same.

smaller than the inequality in nominal wages.¹⁵

Hence the importance of amenities depends on why people move to high price areas. Moretti generalises the Roback model by assuming that the labour supply elasticity is not infinite, i.e. mobility is imperfect.¹⁶ He models the supply of high- and low-skilled workers and their relative wage, which is used as a measure of inequality. The model offers the following testable prediction: Under a demand shock, the number of skilled workers in a city and their nominal wage increase at the same time. If amenities are the dominant explanation, the number of skilled workers (who value amenities relatively more) would increase without a change in their relative wage.

Moretti indeed finds a positive relationship between the college share and college premium which is consistent with the first demand-pull hypothesis. But it is also possible that high-skilled workers were initially attracted by amenities, and firms subsequently located where these workers were and chose technologies which complemented the abundant type of worker. More robust evidence on the demand effect is provided using local labour demand shocks, which are unrelated to amenities, as an instrumental variable for the college share.¹⁷ The results suggest that demand effects are important for explaining the location choices of college-educated workers.

In sum, while amenities certainly play an important role, this evidence suggests that the location of jobs is at least as important in shaping location decisions. In addition, Deckers et al. (2013) find that in Germany higher prices reduce individual life satisfaction, particularly amongst the bottom 40% of the distribution, suggesting that amenities do not fully compensate (also see footnote 4). Therefore, spatial price differ-

¹⁵Ignoring any other other determinants of welfare. This also ignores non-labour incomes and who owns the land.

¹⁶This implies that high- and low-skilled workers do not completely segregate. Of course, one could write down different models of location choice. In Moretti's model, amenities are given exogenously but valued differently by low- and high-skilled workers. Diamond (2013) presents a model where the increasing college share creates amenities endogenously because individuals like having richer neighbours, which is similar to the model by Guerrieri et al. (2013). These amenities are valued more by the high-skilled than the low-skilled. In such a model, amenities are dominant because every extra college worker creates additional amenities. The empirical evidence does not fully support this model. Glaeser and Gottlieb (2008) show that people did not move primarily to high-income areas, as there is no relationship between the income level in 1980 and subsequent population growth.

¹⁷At the national level, Moretti obtains the relative employment growth for high- and low-skilled workers by industry. The local labour demand shock is the weighted sum of these growth rates where the weights are the local employment shares of an industry in the initial year.

ences may have important welfare implications, which is the motivation for studying real wages in this chapter. I extend the existing literature by using a price index which is unique both in its level of spatial disaggregation and coverage of prices of different goods and services.

1.2 The price data

In 2009, a division of the German Federal Ministry for Environment and Building published for the first time a comprehensive regional price index for the entire country (BBSR, 2009). This study built on earlier efforts, for example the study by Krause (1994) which estimated price differentials between East and West Germany. Ströhl (1994) constructed a spatial price index across 50 cities in Germany, but excluded rents.¹⁸ A number of authors extrapolated the Ströhl (1994) data to the rest of Germany econometrically and also incorporated variation in rent or land prices (Roos, 2006; Blien et al., 2009; Kosfeld et al., 2009; Dreger and Kosfeld, 2010). In 2003, the Bavarian state government conducted a more detailed price collection for Bavaria which also included rents (STMWIVT, 2003). Hillringhaus and Peichl (2010) use rent data aggregated into 13 categories for 96 German regions.

The unit of observation in the price data used in this chapter are the 393 district regions (“Kreisregion” in German), which on average have 200,000 residents. Because the compilation of an accurate regional price index requires a large amount of pricing information (7.28M individual prices, 1.40M rents alone), the price data were collected between 2004 and 2008, with most of the price information stemming from between 2006 and 2008. Prices were collected at different locations at the same time. This procedure assumes that prices changed uniformly across regions during this period. I take the average year and conduct my analysis for 2007 which is also the year in which the rental data were collected.

The cross-section price index is calculated as a Laspeyres-type index. The price

¹⁸Earlier, the Federal Statistical Office constructed spatial price indices for 31 cities in 1978 and 4 cities in 1987 (see references in Blien et al. (2009)).

index of district d relative to district r is defined as

$$cspri celevel_d = \frac{\sum_{k=1}^n p_{k,d} \cdot q_k}{\sum_{k=1}^n p_{k,r} \cdot q_k} \quad (1.1)$$

where $q_{k,d} = q_{k,r} = q_k$ since the weights are the same across locations. The variables $p_{k,d}$ and $p_{k,r}$ are the prices of good k in districts d and r respectively. The district of Bonn is chosen as the base category which is similar to earlier spatial price indices for Germany (e.g. Ströhl, 1994).

The consumption weights are held constant across districts at the national consumer price index (CPI) weights, so the index captures pure price effects.¹⁹ According to Hill (1997), the index is thus a multilateral average basket index which has been used in a number of contexts.²⁰ It does not satisfy additivity. This is not a major issue for this analysis because the spatial deflators for specific subgroups of goods are not publicly available. In addition, this particular type of index suffers from the Gerschenkron effect which arises because districts with larger consumption baskets dominate the national weights. This leads to a substitution bias as discussed in more detail in the chapter on global inequality. In the present context, this does not appear too problematic as I would expect the regional variation in consumption baskets across Germany, for example due to different weather, to be relatively small, as argued by Roos (2006). For international price comparisons, the variation in consumption baskets is much greater as discussed in the chapter on global inequality.

The index is based on a bundle of goods and services that covers 72.3% of the value of the bundle that is used for calculating the national consumer price index (BBSR, 2009). This represents an improvement over earlier studies from various countries which cover a smaller share of the consumption basket and typically concentrate on housing costs only.²¹ Housing expenditures make up the largest share of the index

¹⁹The index uses the weighting scheme from the 2005 CPI. In 2013, the German Federal Statistical Office updated these weights but the differences are very small (DESTATIS, 2013).

²⁰For example, it was used by the United Nations Economic Commission for Latin America and Caribbean (Hill, 1997), and by Ströhl's spatial price index for Germany (Roos, 2006). Wingfield et al. (2005) also include it in their study of regional price indices for the UK.

²¹Mogstad et al. (2007) use a municipal-level house price index in Norway. Jolliffe (2006) uses the US Fair Market Rent data and performs a sensitivity test by assuming a positive but small correlation

(30.8%) which is however smaller than for example in the US (42.7%).²² Most of the inter-regional price differentials are driven by differences in housing costs, which is consistent with evidence from other countries.²³ The housing expenditures included in the index are based on new rental agreements.²⁴

Figure 1.1 shows the regional distribution of prices for the entire country. Prices seem to be higher in metropolitan areas and in Southern Germany. The least expensive district, Tirschenreuth in Northeast Bavaria, is 16.6% cheaper than Bonn (the base category) and Munich (city) is 14.4% more expensive. In other words, Munich is 37.2% more expensive than the cheapest district. For the 20 regions of Italy, Pittau et al. (2011) find a very similar differential of 34.1%. Across the 12 regions of the UK in 2004, London is 17.8% more expensive than Wales (Wingfield et al., 2005).²⁵ Across five large US regions, Slesnick (2002) finds that the urban Northeast is 53.8% more expensive than rural households.²⁶ Jolliffe (2006) finds that across the USA,

between rents and the remaining prices. Rienzo (2010) studies spatial price differences in the UK using variation only in rental costs. Moretti (2013) presents an index which includes only housing costs and one which also allows for spatial variation in non-housing costs. As already mentioned, Ströhl (1994) ignores rents for Germany. In their study of 96 German regions, Hillringhaus and Peichl (2010) include only rental data.

²²These expenditure weights include all housing-related expenditures. Including only rents, the shares are 20.3% for Germany and 32.8% for the BLS CPI-U in the US (Moretti, 2013). In Germany, the shares include imputed rents for owner-occupiers.

²³See for example Pittau et al. (2011) for Italy or Moulton (1995) for the USA. Wingfield et al. (2005) confirm this for housing expenditures in the UK, where London is 89.4% more expensive than Northern Ireland in 2004. In contrast, Rienzo (2010) finds that rents are more than twice as expensive in London as in Northern Ireland in 1997.

²⁴It includes new rents for new and existing buildings but does not cover existing rental contracts. The rental data are obtained mostly from newspaper advertisements and online. This may exclude some commercially managed apartments which are marketed through different channels. It is thus assumed that the prices in this housing segment are not systematically different from the rents included in the sample. This approach is different from other studies which use consumption surveys to measure actual rent payments (e.g. Wingfield et al., 2005). Roos (2006) uses observed rents from a 1% sample of households (the German “Mikrozensus”). When using observed rents, it is perhaps harder to control for quality differences, similar to the issues arising in the construction of price indices from unit values in consumption surveys (Deaton and Dupriez, 2011b). Furthermore, the BBSR methodology assumes that the spatial pattern of expenditures by owner-occupiers can be approximated by rents, which follows other studies in this literature (e.g. Jolliffe, 2006; Roos, 2006; Moretti, 2013). For Germany, this assumption is less problematic than in other countries because it has a low proportion of home ownership. Amongst all the OECD countries in 2004, Germany had the second-lowest home ownership rate at 41% (after Switzerland), compared with 71% in the UK and 69% in the USA (Andrews and Sánchez, 2011).

²⁵For 1993, Johnston et al. (1996) finds that London is 16.8% more expensive than Northern Ireland, which is quite similar to Wingfield et al. (2005). Although differences in data sources might make these results not directly comparable.

²⁶This differential seems large. It is important to bear in mind that Slesnick uses unit values from the Consumer Expenditure Surveys, so the results might not be fully comparable (also see footnote 24).

metropolitan areas are 26.6% more expensive than rural areas.

Figure 1.2 shows that the distribution of prices across all districts of Germany is strongly asymmetrical and has a long right-hand tail.²⁷ However, the national distribution hides important differences between East and West Germany. As the figure shows, prices are substantially more concentrated in the East than in the West. The coefficient of variation of the price distribution in the East is approximately half of the value in the West.²⁸ Furthermore, the range of prices in the East lies within the range observed in the West, with the three least expensive districts being in rural Bavaria and Lower Saxony in the West. On average, the price index is 87.5 in East and 92.0 in West Germany, i.e. the West is 5.1% more expensive than the East. No district in East Germany has a price index higher than Bonn, the base category. The ten most expensive districts are all in the West (and, in particular, are all in the three Southern states Baden-Württemberg, Bavaria and Hesse). The large extent of overlap in the price distributions suggests a disaggregated analysis. The level differences in prices between East and West can be partly explained by the fact that goods which are comparatively inexpensive in the East enter the consumer price index only with a low weight (BBSR, 2009).

Hamburg, the most expensive state is 17.6% more expensive than the least expensive state, Saxony-Anhalt (Table 1.1, column 3).²⁹ But most of the variation in prices occurs within states.³⁰ Table 1.1 (column 5) shows that a linear regression of prices on a rural/urban classification and state-dummies explains 60% of the variation

²⁷The price distribution might appear log-normal. However, formal tests of log-normality reject this. From its plotted density, the distribution of the logarithm of the price index appears positively skewed and its tail is somewhat too fat, compared to a normal density (not shown). Studies of the spatial distribution of house prices have found support for log-normality in Tokyo (Ohnishi et al., 2011) but not in a number of UK cities where the house price distribution is closer to a Pareto (MacKay, 2011). If the house price can be written as a multiplicative function of various characteristics, such as in a hedonic regression, a log-normal distribution would not be surprising (Ohnishi et al., 2011).

²⁸The coefficient of variation of the price index is 5.34% (National), 2.62% (East) and 5.35% (West).

²⁹By contrast, using an econometric approximation, Roos (2006) finds that in 2002, the most expensive state (Hesse) is 23.7% more expensive than the least expensive (Mecklenburg-Vorpommern). See footnote 8 for a description of Roos (2006).

³⁰Using a simple analysis of variance, the between-group component explains 30% while the within-group component accounts for 70%. Hence it would not be appropriate to construct a panel of prices using state-level inflation figures. States are the smallest spatial level at which inflation figures are available.

in prices.³¹ Prices increase with urbanisation; core cities in agglomerations for example command a 10pp price premium. After controlling for urbanisation (comparing columns 3 and 5), Southern states still command a price premium, whereas Eastern states are insignificant with the exception of Saxony which is significantly negative and Mecklenburg-Vorpommern which is significantly positive.

1.3 The wage data

I use the SIAB (2010) data, which are based on a 2% random sample of social security records. The database from which the observations are drawn includes employees, benefit recipients and individuals looking for work. In this chapter, I use only the data on employees subject to social security contributions.³² The main advantage of the SIAB database for my purposes is its large sample size.³³ Because the data are based on administrative filings by the employer, they are relatively reliable.³⁴ The downside to being an administrative source is that the data include only individuals covered by the social insurance system, approximately 80% of the workforce, and exclude civil servants, the self-employed and individuals with very low wages (Dustmann et al., 2009).³⁵ They

³¹I use two urbanisation classifications by the BBSR. Column 1 uses the more aggregated classification into four types, ordered by decreasing urbanisation. The base category is rural, so rural districts cost on average 88.53% of Bonn's price level. Core cities cost on average 94.82% of the price level of Bonn. In column 2, I use a more detailed system, which divides districts into nine categories. Column 3 uses only state dummy variables. There is no significant difference between Bavaria and North Rhine-Westphalia, so Munich's price premium (the most expensive district overall), is offset by some of the cheapest districts also being located in Bavaria. Columns 4 and 5 combine state dummies and urbanisation controls.

³²I use the so-called Employee History file (or Beschäftigten-Historik, BeH). Unlike Dustmann et al. (2009), I include apprentices. The results are very similar when apprentices are excluded: The slope coefficients on the regression lines are within one standard error of those reported in Table 1.4 and Figures 1.5 and 1.6. The impact of regional prices on inequality is even somewhat larger than what is reported in Table 1.5.

³³There exist a number of alternative data sources. The income tax data by the Federal Statistics Office (e.g. used by Bach et al. (2009)) do not contain a district identifier due to data protection concerns. The cross-sectional "Mikrozensus" survey contains only limited information on income. The SOEP data are perhaps the most widely used to study questions related to inequality in Germany, particularly because of their detailed information on income and household composition. However, the sample size is too small for my purpose as discussed in footnote 7.

³⁴Employers submit notifications to the social security agency either at the end of the year for ongoing employment or at the termination of the employment relationship.

³⁵Before 1999, individuals in marginal employment (earnings limited to EUR 400 per month in 2007) were not included in the database. Like Dustmann et al. (2009) and Card et al. (2013), I drop them from the data in 2007, since they are very likely to work part-time.

contain no detailed information on hours worked, only whether working full-time or part-time (defined as less than 18 hours per week), so I exclude individuals in part-time employment like other papers using these data (Dustmann et al., 2009; Card et al., 2013). In their analysis of wage inequality, Dustmann et al. (2009) argue that these definitional differences cannot explain the differences in results compared with the SOEP.³⁶ The SIAB data only record earnings from work, so my analysis excludes other sources of income. The unit of observation is the individual, because it is not possible to match records to create households. In summary, the population under investigation includes all individuals in a job covered by social security legislation working full-time, who are living in Germany and have a non-missing place of residence.³⁷ Following Dustmann et al. (2009), each individual is weighted by their duration of employment, i.e. someone who has worked for only half a year receives half the weight of someone who worked the entire year.

During the year, individuals might switch employers which would result in multiple employment spells being recorded. I drop spells with a zero wage, as it is not possible to distinguish them from employment interruptions.³⁸ From the spell data, I construct the average daily (gross) wage in 2007 over the multiple employment spells, weighting each spell by its duration.³⁹ Perhaps the most important drawback of the SIAB data is the right-censoring of the earnings variable. Earnings are only recorded up to an upper earnings threshold in accordance with social security legislation.⁴⁰ The upper earnings

³⁶In contrast to evidence from the SOEP, Dustmann et al. find that wage inequality increased during the 1980s and 1990s. They conclude that these differences are not explained by (1) the exclusion of civil servants and the self-employed, (2) the use of hourly wages in the SOEP, or (3) the omission of annual bonuses and one-time payments from the SOEP.

³⁷Prices correspond to the place of residence, see below. I drop all observations for which the district of residence is missing. I further exclude all individuals living abroad, as I have no price information for these individuals.

³⁸For example, long periods of illness or maternity leave. This affects around 1.7% of spells.

³⁹For each spell, the data record the daily wage and duration (including weekends and holidays). I compute total annual earnings as the sum of total earnings in every spell, and divide by the total duration of all spells within that year. Hence my wage variable is the average daily wage during the time an individual spent working. Blien et al. (2009) follow a different approach. Their daily wage is defined as the wage paid on 30th June, which is the date for which the employer information is recorded. Hence they exclude anyone who is not employed at this date, and they ignore the duration of the spell. Card et al. (2013) take only the longest employment spell, arguing that most full-time workers do not switch employers during a year.

⁴⁰Below that limit, contributions to the pension insurance system are calculated as a percentage of gross earnings. Earnings above the limit are not subject to contributions.

thresholds differ by year and between East and West Germany, accounting for differences in average wages.⁴¹ The daily upper earnings limits in 2007 were EUR 149.59 for East and EUR 172.60 for West Germany.⁴² The distribution of nominal earnings clearly shows spikes at the earnings thresholds (Figure 1.A.1). I apply the East German threshold to Berlin, because it is lower and thus a more conservative approach.⁴³ However, these spikes are not discontinuous jumps, as one would expect from the legislation. This could arise if individuals are employed at the censored wage only for part of the year.⁴⁴ As suggested by Dustmann et al. (2009), I choose a censoring limit slightly below the legislative threshold.⁴⁵ Table 1.A.1 shows the proportion of observations above different thresholds. Overall, the thresholds of EUR 169 in the West and EUR 146 in the East and Berlin seem to be the most appropriate.⁴⁶

The literature that uses these data, does not agree on how to best deal with the censoring issue. I constructed a number of (nominal) earnings variables, of which I report the three most commonly used ones here. The first variable is the raw wage with no adjustments for censoring, which is the approach followed by Baltagi et al. (2009), Riphahn and Schnitzlein (2011), and Schmieder et al. (2012).⁴⁷ The second wage variable simply drops all the censored wages from the analysis. Hence I truncate the wage distribution at EUR 169 in West Germany, and above EUR 146 in East Germany and Berlin. Of course, this would change the population of interest which now becomes

⁴¹In practice, the upper earnings threshold is assigned according to the workplace, i.e. the location of the employer. I match prices and define regions according to the place of residence.

⁴²These are the thresholds for wage and salary earners, which is by far the biggest group in the sample. Miners have a separate pension insurance with a typically higher upper earnings threshold. Because I cannot easily distinguish between the two types of insurance, I apply the threshold for wage and salary earners throughout.

⁴³By law, the respective thresholds are applied to the former eastern and western parts of the city. Unfortunately, it is not possible to distinguish between these two parts in the data, since Berlin now forms only one district and state. The distribution for Berlin shows spikes at the East and West earnings limits. There is also a spike in the East German distribution at the West earnings limit, which could arise because of commuting.

⁴⁴The averaged yearly wage is then less than the censoring limit. Of course, I would observe a higher average wage without censoring. Note that the censoring threshold is applied per spell.

⁴⁵Dustmann et al. find similar results when they reduce the censoring limit by DM 3 or DM 9, i.e. EUR 1.53 or EUR 4.60. I choose an intermediate value of EUR 3.60.

⁴⁶With these thresholds, 8.24% of observations are affected in West Germany, 3.93% in East Germany and 10.43% in Berlin (Table 1.A.1).

⁴⁷The innovation by Riphahn and Schnitzlein is that they censor a constant share of observations in every year. This issue does not arise in my case, as I only consider one year. Schmieder et al. (2012) concentrate on unemployed workers who are affected less by the top-coding.

employees covered by social security legislation with earnings in the middle part of the distribution. This is similar to the main analysis in Dustmann et al. (2009) who focus on the uncensored distribution (up to the 85th percentile). The third wage variable uses a censored normal regression model to impute the censored wages, following the baseline variable in Dustmann et al. (2009).⁴⁸ As in Dustmann et al., the independent variables are three education categories, ten age categories and their interactions.⁴⁹ I allow for heteroskedasticity in the error term for the different age/education groups. I estimate six separate models differentiating by gender and regions (West, East and Berlin). An observation with a wage below the censoring threshold retains its observed wage. The imputed wage for a censored observation is based on the prediction from the model and a draw from the normal error term with the appropriate variance.⁵⁰

Given that a priori there is no preferred way of dealing with the censoring issue, I present results for these three wage variables. For example, the imputed wage relies on the validity of the normality assumption in the error and the predictive power of the wage equation. Like Dustmann, I concentrate on statistics which are less sensitive to the censored part of the wage distribution. For example, I generally use the median instead of the mean and exclude top-sensitive inequality measures such as GE(2). Earlier papers using these data have claimed that their results are not sensitive to how they account for censoring.⁵¹ My results are also robust across the three wage variables

⁴⁸Blien et al. (2009) also impute wages, as suggested by Gartner (2005). Card et al. (2013) adopt a more elaborate imputation model which accounts for within-person heterogeneity by exploiting the panel dimension of the data. Dustmann et al. (2009) select their specification from a number of models with different assumptions about the error term distribution (including a Pareto distribution), and also compare it to uncensored wage data from the German Structure of Earnings Survey. Like Dustmann et al. (2009), I do not weight observations by duration weights because otherwise the regression model does not converge. The results are robust to including duration weights as a control variable.

⁴⁹They define the following three education categories: First, workers with only lower-level secondary education. Second, workers with an apprenticeship or a high school certificate, the Abitur. Third, graduates of a university or polytechnic. The youngest (oldest) age category is defined as younger (older) than 21 (60); the remaining categories correspond to five-year intervals. For each individual, I select the non-missing information recorded in the longest spell.

⁵⁰The number of observations of the imputed wage variable is 0.6% lower than the raw wage (Table 1.2). This is explained by missing values in the independent variables for a censored observation.

⁵¹As already mentioned, Dustmann et al. (2009) find very similar results when using different assumptions on the error term. In their analysis of wage mobility, Riphahn and Schnitzlein (2011) compare the results from the censored data with uncensored data from the SOEP, and find no evidence of a systematic bias. In a private correspondence, Uwe Blien writes that whether or not to impute censored wages does not make a difference to the wage curve analysis in Baltagi et al. (2009).

which deal with the censoring issue in different ways, as described above.

1.4 Constructing real earnings and summary statistics

As mentioned in Section 1.2, the price index is constructed at the level of 393 district regions (“Kreisregion” in German). The earnings data use administrative districts (“Landkreis” and “kreisfreie Stadt”) which are slightly smaller units.⁵² Furthermore, it uses newer boundaries so some districts remain unmatched.⁵³ In sum, I conduct the analysis at the level of 407 districts, ignoring six districts in the Eastern State of Saxony-Anhalt. These observations are excluded throughout, including the analysis of nominal wages and price dispersion.⁵⁴

The real wage is defined as the nominal wage divided by the appropriate price index. I calculate real daily gross earnings of individual i as

$$Realearn_i = \frac{Nomearn_i}{cspricelevel_{D(i)}} \quad (1.2)$$

where $D(i)$ is the district person i lives in. Bonn is the base category, so individuals living there have the same real and nominal wages. Prices refer to the place of residence, not the place of work, which seems appropriate (at the very least for housing expenditures).

Table 1.2 shows summary statistics for the different wage variables, both nominal and real. Note that an analysis of East and West Germany excludes Berlin, which

⁵²The majority of the administrative districts are rural districts (“Landkreis”); larger cities are districts on their own (“kreisfreie Stadt”). The price index is constructed not for these administrative units, but instead small urban districts of less than 100,000 residents are included within the surrounding rural district, resulting in 393 district *regions*. The same price level is assigned to districts in the same district region. The BBSR provides a matching between districts and district regions.

⁵³The price data use 2006 district boundaries, before the number of districts in Saxony and Saxony-Anhalt was reduced. According to the 2006 boundaries there were 439 administrative districts in Germany which was reduced to 413 districts by 2008. The analysis is conducted using the new boundaries and the prices collected for the old district regions are matched with the new district identifiers. I can match districts in Saxony before and after the reform because old districts were simply combined into new units. The situation is more complicated in Saxony-Anhalt, where some old districts were split up. As a result, I cannot match six new districts to prices in this state.

⁵⁴The conclusions are robust to including these districts where possible.

is not part of either region. The imputed wage has slightly fewer observations than the raw wage, because a censored observation might have a missing value for the independent variables in the imputation model. Without adjusting for price differences, the mean (or median) wage is approximately 39% greater in West than East Germany. This differential falls by around 9pp to approximately 30% when taking into account price differentials.⁵⁵ Using SOEP data, Frick (2008) reports mean gross nominal wages per day in 2007 of EUR 92.58 in West and EUR 64.85 in East Germany for individuals working full-time. These are quite close to the figures in Table 1.2 despite all the differences between the two data sources. The data have a very good regional representativeness. On average, there are 1160 observations per district, with a minimum of 201.

1.5 Results at the district level

In the first part of the analysis, I consider the effect of prices on district-level medians and means. This is a natural starting point given that the price deflator acts on a district as a whole. In other words, a person's position in the aggregate wage distribution can only change if accounting for prices shifts an entire district. For comparability with Dustmann et al. (2009), I concentrate on the imputed wage variable but report robustness checks where necessary.

There is almost no overlap in the distributions of median nominal and real wages between East and West Germany. Figure 1.3 shows that for median wages all the Eastern districts are in the bottom 30% in nominal and bottom 25% in real terms.⁵⁶ The overlap even appears to decline after controlling for prices. This lack of overlap in the nominal distribution contrasts sharply with the overlap in the East and West price distributions (Figure 1.2). Table 1.3 shows the ten richest and poorest districts according to median nominal wages together with their price level and their rank in

⁵⁵Using household disposable income (adjusted using an equivalent scale), Goebel et al. (2009) find a smaller nominal premium of 26% which falls by 7pp after the price adjustment.

⁵⁶To create Figure 1.3 I split the national distributions of median district nominal and real wages into 20 ventile groups. The figure shows the regional composition for each of these groups.

the price distribution. The stark contrast between the East and West is evident even from this very coarse illustration: The ten richest districts are all in the West, whereas the ten poorest districts are all in the East. The price ranks in the last column of that table suggest a positive but imperfect (rank) correlation between prices and median wages.⁵⁷

The East-West divide in median wages is also very clear from the maps (Figures 1.4). In the left (middle) panel districts are classified by decile of the distribution of median nominal (real) wages across districts. For example, in the left (middle) panel the lowest category represents districts which are in the bottom 10% of median nominal (real) wages. Even for real earnings (the middle panel) the Eastern districts are still mostly in the bottom two deciles of median wages.⁵⁸ Within West Germany, moving from nominal to real earnings causes quite some re-ranking of districts. The right panel of Figure 1.4 combines the nominal and real maps to illustrate the re-ranking as a result of the price adjustment.⁵⁹ For 262 districts (mostly in West Germany), the median wage falls within one standard deviation of the average median in terms of *both* nominal and real wages (labelled “N;R”).⁶⁰ Most districts in East Germany are below this range in both nominal and real terms (labelled “N-;R-”). Some districts close to Berlin even fall back in real terms given their relatively high price level (labelled “N; R-”). A similar phenomenon albeit at a higher nominal wage is observed for some mostly urban districts within West Germany (labelled “N+;R”). The districts that maintain their top position in nominal and real terms are concentrated in Southern Germany and Lower Saxony.

Figure 1.5 shows the association between earnings and prices in more detail by plotting median nominal (the left panels) and real (the right panels) wages against the price level. The panels in the top row use the entire national sample, while the following two rows repeat for West and East Germany separately. Each panel includes

⁵⁷The Pearson and Spearman rank correlation coefficients between median nominal imputed wages and prices are 0.67 and 0.73 respectively. The third poorest district Rügen has a relatively high price level which might be explained by having a large tourism sector and being an area of natural beauty.

⁵⁸Some coastal Eastern districts even appear relatively poorer in real terms, as they are amongst the poorest districts in nominal terms, but face a relatively high price level (also see Table 1.3).

⁵⁹A similar chart is reported in BBSR (2009) [p. 64], but in terms of GDP per capita.

⁶⁰That is, I take the average and standard deviation of the distribution of nominal (and real) median wages across districts.

a 45° line and a linear fitted line, as well as the coefficient estimates from this linear fit. In the national sample, there is a positive relationship between median nominal wages and prices. As a result, the fitted line for real wages is flatter but remains upward-sloping. For East Germany, the association between nominal wages and prices is weaker compared to the national sample. For the West, the slope of the fitted line for nominal wages is not significantly different from one, i.e. on average there is a one-to-one relationship between nominal median wages and prices. As a result, the fitted line for real wages is flat, i.e. the slope is not significantly different from zero.⁶¹ This suggests a strong equalising effect of prices on median wages in West Germany.

However, this last result is not robust to using the mean instead of the median (Figure 1.6). In West Germany, a positive association between mean real wages and prices remains. In Table 1.4, I summarise the coefficients from separate regressions of mean (and median) nominal (and real) wages on the price index for all wage variables, i.e. the coefficients reported in Figures 1.5 and 1.6. I find that the slope coefficients in real wages are flatter than in nominal wages in all instances, suggesting that prices equalise wage differentials across districts.⁶² This effect appears stronger in West than East Germany. But in almost all specifications, the relationship between nominal wages and prices across districts is not one-to-one, so a significant slope remains for real wages.

In summary, I find that accounting for prices has a bigger effect in West than in East Germany. This is because the distribution of prices is more compressed in the East than the West compared with the distribution of nominal average wages across districts (Figure 1.7).⁶³ In these figures, each of the series is scaled by their respective mean so

⁶¹I labelled a number of interesting districts in the figures for West Germany. Munich (city) is the most expensive district. While it is in the top 10 districts of nominal wages it drops to the bottom 20% of real wages. The Main-Taunus district in the vicinity of Frankfurt/Main is the richest district in nominal terms, and is still in the top 2% of real wages given a high but not excessive price level. Helmstedt, a rural district close to Wolfsburg (the home of Volkswagen) in Northwest Germany, moves up when considering real wages.

⁶²With the truncated wage, the coefficient actually becomes negative in West Germany. Truncating the wage artificially removes some high wages and as a result wages are not allowed to rise sufficiently to offset some of the high price levels.

⁶³The coefficient of variation for prices (median nominal wages) is 5.4% (8.6%) in West and 2.6% (8.3%) in East Germany.

that they all have a mean of one. Consistent with a positive correlation between median wages and prices, the dispersion in real terms is less than the dispersion in nominal terms, and this effect is stronger in West than East Germany. These conclusions are robust to whether mean or median wages are considered.

1.6 Results at the individual level

Table 1.5 shows the effect of accounting for regional price levels on a number of inequality measures, and for the three wage variables. Regional price differences have a small effect on wage inequality. The Gini coefficient changes by approximately 1.5% or 0.4pp which is of a similar order of magnitude for the different wage variables.⁶⁴ The percentage changes in some of the Generalised Entropy (GE) measures and Atkinson (1970) indices are somewhat bigger but still very small. The GE(0) (or Theil-L, or mean log deviation) index is particularly sensitive to income differentials at the bottom of the wage distribution. In the Atkinson index, $A(\varepsilon)$, the weight attached to lower incomes increases with ε , the aversion to inequality. The results from comparing the different GE and Atkinson indices suggest that prices have a smaller effect at the bottom of the wage distribution. These differences are statistically significant using linearised standard errors (in parentheses), with the exception of the A(2) index which attaches a higher weight to incomes at the bottom.

The district-level analysis (Section 1.5) suggests that there are important differences in the price effects between East and West Germany, so I report separate results for these regions. Accounting for prices has a weaker effect on inequality in the two regions compared with the aggregate sample. Consistent with the district-level results, the effect is weaker in East than West Germany. While for West Germany the price effect remains statistically significant (with the exception of the A(2) index), this is not the case in East Germany.⁶⁵ With the exception of the imputed wage measure, inequality

⁶⁴Using the SOEP data, Goebel et al. (2009) also find a decline in the Gini index of 0.4pp. This contrasts with the experience in other European countries, such as Italy where the Gini index of disposable income falls by 4.5% after accounting for price differences at the level of 20 regions (Pittau et al., 2011).

⁶⁵For example, the Gini index of the imputed wage measure in East Germany declines by 0.0013

in the East is higher than in the West which is the opposite of disposable income in the SOEP data.⁶⁶

The bottom panel of Table 1.5 decomposes the $GE(0)$ measure into the between- and within-components by district.⁶⁷ The price adjustment affects only the between-district component of the decomposition, as prices are constant within a district. Between-district inequality in the national sample falls by 40% or more depending on the wage variable. Consistent with the above evidence, the price adjustment has a greater effect on the between-component in West than East Germany. But this component only accounts for approximately 5% of total inequality of nominal wages in the national sample. Hence the overwhelming majority of wage inequality in Germany exists within rather than between districts.⁶⁸ Because the price adjustment only affects between-district inequality, this explains why the effect on aggregate inequality is small (although statistically significant).

after the price adjustment. Note that standard errors decrease by having a large sample size, so part of this effect could arise from the Eastern sample being much smaller.

⁶⁶The distribution of monthly disposable income amongst individuals in private households has a Gini of 0.272 for all Germany, and 0.276 and 0.234 for West and East respectively (SOEP Group, 2014). Using annual incomes (which is closer to my specification) and including imputed rents, increases the Gini to 0.286 nationally, 0.290 in West and 0.239 in East. The two datasets use quite different populations (employees liable to pay social insurance), units of analysis (individuals v. households), and welfare measures (censored wages v. disposable income), so it is not surprising that the results are different.

⁶⁷I only report the $GE(0)$ measure for two main reasons. First, the Gini index is not subgroup-decomposable. While the Atkinson index is decomposable, it cannot be expressed as a weighted average of the within- and between-group inequalities (Bourguignon, 1979). Second, I choose the $GE(0)$ over other generalised entropy measures because its weights in the decomposition are the population shares. Population shares are not affected by the price adjustment. For the other GE measures, income shares enter the decomposition.

⁶⁸For comparison, in Italy the between-group share of the $GE(1)$ measure falls from 14.4% to 8.1% after accounting for prices at the level of 20 Italian regions (Pittau et al., 2011). There is some evidence to suggest that between inequality is more important in poorer countries. Using a $GE(0)$ inequality measure and somewhat smaller subgroups than in my analysis, the between-share is around 5-10% in Ecuador, compared with 15-20% in Madagascar and Mozambique (Elbers et al., 2005). However, Förster et al. (2005) find a small between-share for the Czech Republic, Hungary and Poland (e.g. 1.3% in Poland in 1992), which is even smaller than what they find in Italy. Another strand of the literature examines whether the dispersion in per capita GDP across regions varies with economic development (Lessmann, 2014). These studies only capture the between-region level and not the share.

1.7 Conclusion


This chapter exploits a new price index which for the first time allows an accurate comparison of price levels across the districts of Germany. I combine the price data with wage information from social security records, which provide a sufficient number of observations to perform a very disaggregated analysis. I use a number of methods to address the top-censoring of the wage variable arising from social security legislation.

I find that accounting for prices reduces the wage differentials between districts. I find that on average, median nominal wages (for the preferred wage measure) move one-to-one with the price level in West Germany, suggesting a strong equalising role of prices. By contrast, this effect is weaker in East Germany, where the distribution of prices is substantially more concentrated than the distribution of nominal average wages across districts. Furthermore, I find that accounting for prices reduces the West-East differential by some 9pp, but even in real terms the average wage in West Germany is approximately 30% greater than in the East. A possible interpretation is that nominal wages in East Germany are held down by low labour productivity (Fuchs-Schündeln and Izem, 2012).


Accounting for prices reduces between-district inequality by about 45% in Germany. This effect is weaker in East (around 35%) but stronger in West Germany (almost 60%). Aggregate inequality in real wages is lower than in nominal wages but the difference, although statistically significant, is small. The Gini index decreases by approximately 0.4pp. The small effect on overall inequality is explained by the fact that most of the inequality in wages is within, rather than between districts. This is quite remarkable given that I consider very small regional units. Accounting for cross-district price differences only affects between-district inequality.

1.8 Main Tables and Figures

Figure 1.1: Regional price index (BBSR, 2009)



Regional price index (Bonn=100)



Source: BBSR (2009), p. 60

Figure 1.2: Density of price index




Figure 1.3: Regional composition of the distribution of median wages




Table 1.1: Regression of the price index on rural/urban and state indicators

Price index	(1)	(2)	(3)	(4)	(5)
Core city	6.287***			6.986***	
High density periphery	3.777***			3.975***	
Rural periphery	0.0608			0.685	
Agglomerations: Core cities		8.679***			10.65***
Agglomerations: High density		8.181***			9.135***
Agglomerations: Medium density		5.737***			7.039***
Agglomerations: Rural		2.374*			4.204***
Urbanised regions: Core cities		5.845***			6.654***
Urbanised regions: Medium density		3.134***			3.423***
Urbanised regions: Rural		0.975			1.324
Rural: Major density		2.050*			1.760*
Schleswig-Holstein (West)			0.333	3.075**	5.009***
Hamburg (West)			9.280*	7.429*	7.008*
Lower Saxony (West)			-3.433***	-0.994	1.407
Bremen (West)			-1.170	-3.021	-1.445
Hesse (West)			2.830**	4.223***	5.123***
Rhineland-Palatinate (West)			-1.590	0.0587	2.440**
Baden-Württemberg (West)			3.368***	4.654***	6.566***
Bavaria (West)			-0.369	3.035***	5.010***
Saarland (West)			-2.804	-2.145	-3.116*
Berlin			1.080	-0.771	-1.192
Brandenburg (East)			-4.070***	-0.206	0.788
Mecklenburg-Vorpommern (East)			-3.581**	1.090	3.984***
Saxony (East)			-5.283***	-3.953***	-2.329*
Saxony-Anhalt (East)			-5.895***	-2.849*	0.156
Thuringia (East)			-5.107***	-1.753	0.848
Constant	88.53***	87.30***	92.12***	86.99***	83.74***
R-squared	0.244	0.348	0.297	0.484	0.605
adjusted R-squared	0.239	0.334	0.270	0.460	0.582
Number of observations	407	407	407	407	407

Dependent variable: price index (100=Bonn; range from 83.4 to 114.4)

(1): Omitted category is Rural; (2): Omitted category is Rural: Minor density

(3): Omitted category is North Rhine-Westphalia; (4): Omitted categories are Rural and North Rhine-Westphalia;

(5): Omitted categories are Rural: Low population density and North Rhine-Westphalia.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.2: Summary statistics for different wage variables in 2007, nominal and real

Wage variable	National			West			East		
	Mean	P50	N	Mean	P50	N	Mean	P50	N
Nominal raw wage	86.1	81.7	472,096	90.2	86.0	379,608	66.1	60.1	75,410
Nominal wage (truncated at 169, 146)	78.3	77.5	436,067	82.0	81.6	348,325	62.1	58.5	72,446
Nominal wage (imputed above 169, 146)	91.0	81.3	469,126	95.6	85.6	377,127	67.7	59.9	75,168
Real raw wage	92.4	88.2	472,096	95.9	92.2	379,608	75.5	68.9	75,410
Real wage (truncated at 169, 146)	84.4	83.7	436,067	87.5	87.4	348,325	70.9	67.0	72,446
Real wage (imputed above 169, 146)	97.6	87.9	469,126	101.6	91.8	377,127	77.3	68.7	75,168

All statistics are duration-weighted. Real wages at Bonn's prices.

Figure 1.4: Regional distribution of nominal and real median wages per day (by decile); wages imputed above 169, 146




Figure 1.5: Median wages and prices across districts: National, West and East




Figure 1.6: Mean wages and prices across districts: National, West and East




Table 1.3: The 10 richest and poorest districts and their prices


District region	State	Region	P50 nom. daily wage	Price index	Price rank
Main-Taunus-Kreis	Hesse	West	110.38	104.4	9
Munich (rural)	Bavaria	West	108.37	109.6	2
Wolfsburg	Lower Saxony	West	107.39	91.2	160
Hochtaunuskreis	Hesse	West	106.65	105.4	5
Böblingen	Baden-Württ.	West	104.5	101.3	21
Eichstätt	Bavaria	West	100.07	91.6	151
Helmstedt	Lower Saxony	West	99.65	87.6	310
Ebersberg	Bavaria	West	99.42	104	10
Munich (city)	Bavaria	West	98.8	114.4	1
Fürstentfeldbruck	Bavaria	West	98.15	104.7	7
...	...	...	...	...	...
Kyffhäuserkreis	Thuringia	East	55.23	85.7	381
Elbe-Elster	Brandenburg	East	55.07	84.7	397
Demmin	Meck.-Vorp.	East	54.91	85.6	383
Altenburger Land	Thuringia	East	54.8	86.9	339
Görlitz	Saxony	East	54.27	84.4	402
Ostvorpommern	Meck.-Vorp.	East	53.59	90.3	193
Erzgebirgskreis	Saxony	East	53.34	85.4	386
Müritzt	Meck.-Vorp.	East	52.53	86.6	351
Rügen	Meck.-Vorp.	East	49.58	90.7	172
Uecker-Randow	Meck.-Vorp.	East	49.2	86	374
Nominal wages are imputed above EUR146 or EUR169 per day.					
Districts are ranked highest to lowest by their median nominal wage.					

Table 1.4: Relationship between median and mean wages and prices across districts


	Raw wage		Truncated wage		Imputed wage	
	nominal	real	nominal	real	nominal	real
Using the median wage at the district level						
National	1.654*** (0.089)	0.846*** (0.100)	1.204*** (0.083)	0.424*** (0.093)	1.601*** (0.089)	0.796*** (0.100)
West	1.062*** (0.060)	0.158** (0.065)	0.619*** (0.049)	-0.249*** (0.053)	1.005*** (0.059)	0.104 (0.064)
East	1.323*** (0.200)	0.697*** (0.226)	0.980*** (0.174)	0.340* (0.197)	1.296*** (0.198)	0.668*** (0.224)
Using the mean wage at the district level						
National	1.698*** (0.080)	0.852*** (0.090)	1.146*** (0.070)	0.351*** (0.079)	2.141*** (0.093)	1.263*** (0.105)
West	1.183*** (0.057)	0.253*** (0.061)	0.658*** (0.042)	-0.210*** (0.046)	1.587*** (0.073)	0.614*** (0.079)
East	1.316*** (0.196)	0.634*** (0.222)	0.904*** (0.148)	0.225 (0.168)	1.494*** (0.221)	0.811*** (0.251)
The table reports the coefficient of a regression of nominal (or real) district mean (or median) wages on the price index. Standard errors in parentheses.						

Figure 1.7: Comparing distribution of prices, median and mean wages by region


(a) West, median



(b) East, median



(c) West, mean



(d) East, mean




Table 1.5: Effect of prices on inequality in 2007 by region

	Raw wage			Truncated wage			Imputed wage		
	nominal	real	%diff	nominal	real	%diff	nominal	real	%diff
A. National									
Gini	0.2895 (0.0003)	0.2851 (0.0003)	-1.54	0.2710 (0.0003)	0.2678 (0.0003)	-1.18	0.3295 (0.0005)	0.3242 (0.0004)	-1.59
GE(0)	0.1675 (0.0004)	0.1632 (0.0004)	-2.54	0.1512 (0.0004)	0.1486 (0.0004)	-1.77	0.2056 (0.0005)	0.2002 (0.0005)	-2.63
GE(1)	0.1363 (0.0003)	0.1324 (0.0003)	-2.86	0.1212 (0.0003)	0.1187 (0.0003)	-2.00	0.1895 (0.0006)	0.1837 (0.0006)	-3.08
A(2)	0.3783 (0.0139)	0.3720 (0.0137)	-1.66	0.3573 (0.0148)	0.3528 (0.0145)	-1.26	0.4113 (0.0133)	0.4047 (0.0131)	-1.61
A(1)	0.1542 (0.0003)	0.1506 (0.0003)	-2.34	0.1404 (0.0003)	0.1381 (0.0003)	-1.64	0.1858 (0.0004)	0.1814 (0.0004)	-2.37
A(0.5)	0.0723 (0.0001)	0.0704 (0.0001)	-2.63	0.0650 (0.0001)	0.0638 (0.0001)	-1.83	0.0930 (0.0003)	0.0905 (0.0003)	-2.74
N	472,096	472,096		436,067	436,067		469,126	469,126	
B. West Germany									
Gini	0.2798 (0.0003)	0.2774 (0.0003)	-0.85	0.2622 (0.0003)	0.2612 (0.0003)	-0.36	0.3220 (0.0005)	0.3186 (0.0005)	-1.05
GE(0)	0.1583 (0.0004)	0.1561 (0.0004)	-1.38	0.1433 (0.0004)	0.1427 (0.0004)	-0.44	0.1970 (0.0006)	0.1938 (0.0006)	-1.63
GE(1)	0.1282 (0.0003)	0.1261 (0.0003)	-1.64	0.1144 (0.0003)	0.1138 (0.0003)	-0.54	0.1810 (0.0006)	0.1771 (0.0006)	-2.13
A(2)	0.3711 (0.0184)	0.3672 (0.0179)	-1.04	0.3515 (0.0196)	0.3494 (0.0190)	-0.59	0.4057 (0.0176)	0.4014 (0.0171)	-1.05
A(1)	0.1464 (0.0003)	0.1446 (0.0003)	-1.27	0.1335 (0.0003)	0.1330 (0.0003)	-0.41	0.1788 (0.0005)	0.1762 (77.0000)	-1.48
A(0.5)	0.0684 (0.0002)	0.0673 (0.0002)	-1.47	0.0616 (0.0002)	0.0613 (0.0002)	-0.48	0.0892 (0.0003)	0.0876 (0.0003)	-1.81
N	379,608	379,608		348,325	348,325		377,127	377,127	
C. East Germany									
Gini	0.2959 (0.0007)	0.2947 (0.0007)	-0.40	0.2765 (0.0007)	0.2756 (0.0007)	-0.31	0.3138 (0.0010)	0.3126 (0.0010)	-0.40
GE(0)	0.1721 (0.0010)	0.1712 (0.0010)	-0.57	0.1557 (0.0010)	0.1551 (0.0010)	-0.40	0.1900 (0.0013)	0.1889 (0.0013)	-0.59
GE(1)	0.1431 (0.0007)	0.1420 (0.0007)	-0.76	0.1258 (0.0007)	0.1251 (0.0007)	-0.54	0.1699 (0.0013)	0.1685 (0.0013)	-0.78
A(2)	0.3648 (0.0044)	0.3639 (0.0045)	-0.25	0.3452 (0.0046)	0.3447 (0.0046)	-0.16	0.3800 (0.0044)	0.3790 (0.0044)	-0.26
A(1)	0.1581 (0.0008)	0.1573 (0.0008)	-0.53	0.1442 (0.0008)	0.1437 (0.0008)	-0.37	0.1731 (0.0010)	0.1722 (0.0010)	-0.53
A(0.5)	0.0749 (0.0004)	0.0744 (0.0004)	-0.65	0.0670 (0.0004)	0.0667 (0.0004)	-0.46	0.0849 (0.0006)	0.0844 (0.0006)	-0.66
N	75,410	75,410		72,446	72,446		75,168	75,168	
D. Subgroup decomposition by district for GE(0): National, West and East Germany									
N-With.	0.1583	0.1583	0.00	0.1445	0.1445	0.00	0.1934	0.1934	0.00
N-Betw.	0.0092	0.0050	-46.18	0.0067	0.0040	-39.85	0.0122	0.0068	-44.44
W-With.	0.1548	0.1548	0.00	0.1417	0.1417	0.00	0.1916	0.1916	0.00
W-Betw.	0.0035	0.0013	-62.24	0.0017	0.0010	-37.97	0.0054	0.0022	-59.11
E-With.	0.1694	0.1694	0.00	0.1542	0.1542	0.00	0.1868	0.1868	0.00
E-Betw.	0.0027	0.0017	-36.36	0.0016	0.0010	-38.82	0.0032	0.0021	-34.61

Notes: Employment-duration weights are used throughout. Berlin is excluded from West and East.

Linearised standard errors in parentheses.

1.A Appendix Tables and Figures

Figure 1.A.1: Nominal earnings for Berlin, East and West Germany: Raw wages




Table 1.A.1: Number and proportion of censored observations in 2007

Threshold (EURO)	East		West		Berlin	
	No. obs.	Perc.	No. obs.	Perc.	No. obs.	Perc.
172.60			25,045	6.60%	572	3.35%
169.00			<i>31,283</i>	<i>8.24%</i>	726	4.25%
150.00			46,389	12.22%	1,217	7.13%
149.59	965	1.28%			1,254	7.34%
146.00	<i>2,964</i>	<i>3.93%</i>			<i>1,782</i>	<i>10.43%</i>
130.00	4,625	6.13%			2,480	14.52%
Total obs.	75,410		379,608		17,078	

Notes: Table shows the number and proportion of observations above a given threshold income. Bold numbers are the official thresholds. Numbers in italics refer to the thresholds used in the estimation.

Chapter 2

Top Incomes in the USA, 1960-2005:

Family Size and Factor Income

Composition

Introduction

Over the last decade or so, a substantial body of literature has documented the increase in income inequality at the very top, typically measured as the share of total income received by a top fractile, such as the top 1%. This literature has used tax records, which provide more accurate data about the very top tail. In the US and some other mostly Anglo-Saxon countries, top shares continued to rise substantially, whereas inequality in the rest of the distribution measured from household surveys increased much slower. This research has received a lot of attention in the popular media, not least in the form of the “We are the 99%” slogan of the Occupy movement. Although the top 1% are a small number of people, they command a substantial share of resources (almost 20% of total income in the USA in the latest years), pay an even larger share of taxes in a progressive tax system, and as a result of their economic resources perhaps have a disproportionate amount of political influence (Solow, 2014).

Using data based on Internal Revenue Service (IRS) filings, this chapter addresses

two challenges that have been made to the literature on top income shares in the USA. First, this literature uses tax units as its unit of analysis and ignores the fact that these could be a single person or a married couple, living with a number of dependents. I investigate whether the conclusions about rising top income shares stand up once I allow for differences in tax unit size, such as differences between singles and couples. Between the 1960s and 2000s, the average tax unit size declined from 2.6 to almost 2 persons. Over the same period, the proportion of tax units who are married declined by 10pp (from 60% to 50%) amongst all tax units, but by only 5pp (from 90% to 85%) for the top 10% (Saez, 2004). Such a demographic shift would have increased top income shares by itself as single tax units, who are poorer, are added faster at the bottom than the top. Of course, the effect on the estimated top income share depends on whether this decline was spread evenly along the income distribution. This issue is particularly important in the USA, because in other countries, such as Canada, individual incomes are separated in the tax data. When I account for tax unit size in this chapter, top income shares decrease, but the two series track each other very closely.

Second, I examine the claim by Piketty and Saez (2007b) that high earners have replaced the rentiers at the top of the income distribution. At the same time as the share of income going to the top 1% increased, their share of income from labour also rose substantially (Piketty and Saez, 2007b). This increase in the labour share amongst the top could arise from a change in the association between labour and capital (i.e. an increasing share of the top 1% receive income from both sources), or because the share of tax units in the top who receive only capital income declines while there remain separate classes of (high-paid) labourers and capitalists. Piketty and Saez tend towards the latter interpretation and consider the distribution in class terms, i.e. rentiers and workers. However, class relations have become more complicated, as individuals have both sources of income. I find that the association between capital and labour incomes has indeed increased, at least at the top. Furthermore, the association is asymmetric: Those tax units at the top of the labour distribution are more likely to also be at the top of the capital distribution, compared with tax units at the top of the capital

distribution being at the top of the labour distribution.

This chapter builds on a series of papers by Piketty and Saez, in particular Piketty and Saez (2007b), which is slightly different from the earlier papers.¹ My analysis departs from Piketty and Saez (2007b) in two ways. First, they use tabulations of tax returns whereas my analysis relies on micro data, and therefore I can only start in 1960.² Second, and as a result of using tabulations, they rank tax units by Adjusted Gross Income (AGI), while I use market income (the differences are explained in Section 2.2.2). Some of the analytical work in Part II (particularly Section 2.5) is my contribution to a working paper comparing the US and Norway (Aaberge et al., 2013). In these two countries the tax data capture capital gains more completely than in other countries.³

This chapter is structured in six sections. Section 2.1 reviews the literature and motivates the analysis. In Section 2.2, I describe the data and methodology, especially how the unit of analysis and the income concept are defined in the tax data. Part I of the analysis includes Section 2.3 which discusses the results of adjusting for tax unit size, measured as the number of exemptions claimed on a tax return. As a first step, this section also includes a replication of Piketty and Saez (2007b). Part II of the analysis which comprises Sections 2.4 and 2.5 focusses on the association between capital and labour incomes. In Section 2.4, I present results of a factor income decomposition applied to top income shares. The measure of association in this decomposition is not independent of changes in the marginal distributions. Hence, Section 2.5 uses a framework based on the copula function to derive a rank-based measure of association which does not suffer from this limitation. This is done at the tax unit level without a

¹The 2007b paper differs from the 2001 and 2003 papers in two aspects: First, the earlier papers included social insurance and unemployment benefits, which are dropped in the later paper in order to obtain a consistent estimate of market income. Second, the papers differ in the construction of the control total, which is defined as the total income of the population, i.e. the denominator in the top income share. Piketty and Saez (2001) only included the capital gains received by the top 10%, whereas Piketty and Saez (2007b) use all reported capital gains. It is unclear how the 2003 paper treats capital gains.

²Piketty and Saez (2007a), who measure the progressivity of the US tax system, use micro data.

³For example, in the UK capital gains are taxed separately, and therefore not reported in the income tax data. In Germany, capital gains were largely tax-exempt if they were realised after a certain time limit, e.g. ten years for real estate and one year for stocks and shares.

size adjustment, given that the time trend of top income shares is robust to adjusting for tax unit size (as shown in Section 2.3) and in the interest of comparability with Piketty and Saez (2007b). Section 2.6 concludes. In the Appendix, I address several criticisms of the evidence from tax records, including tax evasion, changes in the tax code, and issues related to the income concept, in particular the treatment of capital gains and the use of before-tax incomes.

2.1 Literature review and motivation

In the US, the share of total income going to the top 1% of income recipients has approximately doubled over the last forty years from less than 10% in the 1960s to almost 20% in the 2000s (Figure 2.1 based on Saez, 2013). Put differently, incomes at the top have grown faster than the average. The top income shares observed nowadays are similar to the early part of the 20th century, giving rise to a U-shaped time series.⁴ Top incomes suffered during the World Wars and the Great Depression and only started to recover in the 1980s. In proportional terms, the income share of the top decile excluding the top 1% (i.e. P90-P99) has increased less since World War II (WWII), while those fractiles inside the top 1%, e.g. the top 0.01%, have seen an even stronger rise (Piketty and Saez, 2007b). Similarly, evidence from household surveys suggests that inequality increased slower in the part of the distribution that excludes the top.⁵

Over the last decade, a series of papers have provided evidence on top income shares from many other countries.⁶ As pointed out by Atkinson et al. (2011), there

⁴In contrast, the Kuznets' (1955) hypothesis says that income inequality follows an *inverse* U-shape as the country develops.

⁵Using the Current Population Survey (CPS), Gottschalk and Danziger (2005) show a strong increase in family income inequality (as measured by the P90/P10 ratio) in the late-1970s and primarily the 1980s, a slow-down in the 1990s and a levelling off in the early 2000s. In contrast, Figure 2.1 shows that the top shares continued to rise in the 1990s and the mid-2000s. Burkhauser et al. (2012) show that some of these differences can be reconciled once the income concept (gross income v. post-tax post-transfer income) and the unit of analysis (tax units v. households) are defined consistently, and when they use internal-CPS data which suffer less from top-coding. However, while this is true for the P90-P99 group, CPS estimates of the top 1% share are lower than in the tax data because of top-coding and underestimation of capital gains and stock options, which are concentrated at the top (Atkinson et al., 2011).

⁶These papers are collected in the two edited volumes by Atkinson and Piketty (2007; 2010). The data on top income shares are available as the World Top Incomes Database (WTID) (Alvaredo et al., 2013b). For the US, the original papers include Feenberg and Poterba (1993; 2000) and Piketty and

appears to be a contrast between the Anglo-Saxon countries on the one hand, where over the 20th century top shares follow the said U-shaped pattern, and Central Europe and Japan on the other hand, where top shares did not increase.⁷ In the Nordic and Southern European countries, top income shares increased in the late 20th and early 21st century, but not enough as to make up for the losses in the earlier part of the 20th century. Among the developing countries, there exists substantial heterogeneity but many point towards increasing top income shares since the 1980s.

During the same period as their income shares increased, the composition of top incomes changed dramatically. Figure 2.2 shows that over this 100-year period the share of income of the top 1% from capital declined, while the labour share increased.⁸ The series excluding capital gains shows this most clearly, but if one ignores some of the spikes in the mid-1980s and 2000s this trend is also visible when I include capital gains. Piketty and Saez (2003) argue that the Great Depression and the World Wars had lasting effects on top capital incomes, which were not able to recover in the post-WWII period, possibly due to progressive income taxation limiting capital accumulation. However, capital incomes remain an important income source at the top, which is even more obvious using a more comprehensive measure of capital incomes (Wolff and Zacharias, 2009). The capital share (including capital gains) (Figure 2.2) has certainly declined from around 50-60% in the early 20th century, but at around 30% in the 2000s it is still substantial.

Hence the “working rich” and the “coupon-clipping rentiers” “now appear to *co-habitate* the top end of the income distribution” [emphasis in the original] (Wolff and Zacharias, 2009, p. 108). As a result, the Ricardian model of rich landlords, poor labourers and capitalists somewhere in between, has been replaced over these 100 years

Saez (2001; 2003).

⁷However, comparing the level rather than the change, top shares in some of these countries, such as Germany, were higher than the US or UK between the end of WWII and the mid-late 1980s.

⁸In this chapter, income from S-corporations is included in capital income, which will be explained in more detail on p. 94. In the data underlying Figure 2.2, S-corporations are included in entrepreneurship income, which in turn is allocated 2/3 to labour and 1/3 to capital. In Figure 2.2, fractiles are defined according to income excluding capital gains, because this is the only income composition series presented in Saez (2013).

by an economy in which the sources of income have become more mixed.⁹ However, large estates might make a comeback because the new working rich can accumulate wealth in an environment where top tax rates are lower than in the 1960s (Piketty and Saez, 2007a).¹⁰ This change in the composition of top incomes is not observed in all countries. In Sweden, the share of capital at the top of the distribution is back to 1945 levels after having fallen until the 1970s (Roine and Waldenström, 2008). Similarly, in Finland the rise in top shares since the mid-1990s was mostly driven by increases in capital incomes (Jäntti et al., 2010). In the UK, the increase in the top income share in the decade following 1998, however, was driven predominantly by increased employment income, particularly in the financial industry (Bell and Van Reenen, 2013), similarly to the US.

2.2 The data and methodology

I use the Statistics of Income Public Use Tax Files by the US IRS over the period 1960 to 2005, which are (nearly) annual public-use samples of the federal income tax returns.¹¹ There exist no micro-files for 1961, 1963 and 1965, and the files available in the first half of the 1960s contain only limited information and are not very robust.¹² I use the data up to 2005, which is the period for which Emmanuel Saez has constructed

⁹In a Ricardian class model, capitalists are on top of the capital distribution and bottom in the wage distribution, i.e. the correlation of labour and capital incomes is negative. Many years ago, Lydall wrote that “it is well known that many landowners are poor, many employees earn more than some capitalists, many property-owners work and many workers own property” (Lydall, 1968, p. 2). And Atkinson argues that “pure ‘rentiers’ may have virtually disappeared” (Atkinson, 2009, p. 6).

¹⁰In the US, the top marginal tax rate declined from 91% in early 1960s to 28% in 1986. By 1993, the top rate had been increased to 39.6% and in 2003 it was cut to 35%. For a more detailed discussion see Section 2.A.2.

¹¹Throughout this chapter, I use these micro data. In the estimation of historic top shares additional methodological issues arise (Atkinson et al., 2011). The historic data are typically in the form of tax record tabulations by income ranges. In order to derive the distribution, an interpolation technique has to be used. Furthermore, historically only a small fraction of the population paid income taxes. As a result, for these historic analyses, the control total (the denominator of the income share) has to be based on national accounts.

¹²The micro-files in 1960, 1962 and 1964 contain fewer tax return variables than subsequent years. This is especially true for the 1960 file. As a result, Piketty and Saez (2007b) use the tabulated data for the first half of the 1960s. In my analysis, however, I need to use micro data and I have chosen to include those years. However, one should interpret the results in the first half of the 1960s with caution.

variables which account for changes in the tax code.¹³ I use these variables with some minor corrections. The data are a random sample of tax records stratified by income and filing status in order to ensure that high-income tax units are well represented. In order to minimise the risk of disclosure, individual identifiers are removed, entries are rounded to four significant figures, for the high-income individuals some fields are blurred, and a very small number of records are removed.¹⁴ Despite these adjustments, these data are the best available source of information for high incomes.

The files contain returns which are filed during a particular calendar year.¹⁵ The vast majority of returns refer to the previous calendar year; e.g. most of the returns filed during 2006 refer to tax year 2005 and hence this is called the 2005 Income Tax file. A small number of returns refer to an earlier tax year and, by the same logic, some returns for the tax year in question will be filed in future years. Because it only influences a small number of observations this is unlikely to affect my results.¹⁶

2.2.1 Unit of analysis

The unit of analysis is a tax unit as defined under US tax law.¹⁷ When I adjust incomes for tax unit size using equivalence scales, I also present results at the individual level. I include all tax units who file a return, including those that do not claim any taxpayer

¹³I also had access to the non-harmonised data for 2006, but nothing later.

¹⁴In the 2006 file 22 records were removed for disclosure avoidance. For a full description of the sampling procedure see the General Description Booklets and Winglee et al. (2002). The disclosure avoidance procedures have become more stringent over time. The technique of multivariate blurring applied since the 1980s essentially involves averaging over tax units sorted by the particular income sources which are deemed sensitive. The differences between the full and the public use samples in terms of weighted sums for the various entries are small (e.g. GDB, 2006).

¹⁵Taxpayers can also choose a tax year which does not coincide with the calendar year (subject to IRS approval). The data contain information on the month in which the taxpayer's tax year ended only in some years. In 2005, for 99.99% of records the tax and calendar years coincide.

¹⁶In the 2006 file, 96.93% of returns refer to the tax year 2006; 2.11% to the 2005 tax year; 0.71% to the 2004 tax year and 0.26% to the 2003 tax year. Shifting returns across tax files is difficult since the sampling procedures, variable definitions etc. might change across files. For example, the sampling weights are designed to return population-level estimates within every Income Tax file, but are not transferable across years. Furthermore, it could be argued that the characteristics of those returns, which are due but have not yet been filed, are similar to the returns which have been filed for previous years (IRS, 2010).

¹⁷In 2000, US tax law recognised five different statuses: single, married filing a joint return, married filing separately, head of household (unmarried, paying more than half of the upkeep of a home, and in most cases having a qualifying dependent), and surviving spouse (widow(er) with dependent child). A different filing status attracts a different tax liability.

exemption.¹⁸ Tax units are typically smaller than households, which are defined by a common residence and sharing of resources. Note that “families” are distinct from households and are often defined by birth/marriage/adoption.¹⁹ For example, two flatmates could be counted as one household but they would be taxed separately. On the other hand, some individuals, such as dependent parents or children living away from home, might be counted as a separate household but are part of the same tax unit. Between 1970 and 2005, the number of households (reported by the US Census Bureau) is between 19.1% and 12.4% smaller than the number of tax units. Using the 2005 PSID, Hungerford (2010) estimates 75% of households to have one tax unit, and another 17% to have two tax units. The unit of analysis might explain some of the differences between the analyses based on the IRS and Current Population Survey (CPS) data (see footnote 5).²⁰

Sampling weights are used to account for the stratification (by income and filing status). This ensures that the sample is representative of all the tax returns filed in a particular year. Not every tax unit files a tax return; historically there were high exemption levels and income taxes only applied to the most affluent taxpayers.²¹ Parts

¹⁸In 2005, 7.4% of returns do not claim any primary or secondary taxpayer exemptions, i.e. they claim to be dependent on another taxpayer. 0.1% of these taxpayers are in the top 10% of income recipients (incl. capital gains), with even smaller proportions in the top 5%, top 1% and top 0.5% (see panel C of Table 2.1 for their average incomes). Before 1984 (with the exception of some very small proportions in 1962 and 1964), there are no tax units without any taxpayer exemptions. The dependent tax units could be children aged 14-18 or students aged 19-24, but the data have no demographic information. In the US, children have to file a return if they have earned income above a certain limit. If they receive only capital income below a limit, parents can choose to include them on their return. I have decided to include these returns, because this seems better than the following alternatives. First, I could drop them and treat them like a non-filer. However, they are not non-filers because they have a positive taxable income which should be included. Furthermore, by treating them as non-filers, I would effectively assume that they are at the bottom of the distribution. Second, I could drop them but include their income in the control total. However, this would be equivalent to assuming that none of these observations is in the top fractiles, which I think would be an overly strong assumption. It is true that my approach mis-measures household income, because their income should be matched with whoever they are dependent on. But this critique applies more widely as I we cannot match tax units to create households.

¹⁹Note that Piketty and Saez (2007b) refer to tax units as “families” (as opposed to “households”). I do not adopt this terminology here to avoid any confusion, because other authors use families and households interchangeably. Furthermore, according to the US Census Bureau, families are a subset of households: Families consist only of people related by birth/marriage/adoption (Lofquist et al., 2012). Hence a household composed of non-relatives, such as unmarried partners, would not be counted as a family.

²⁰For example, in the mid-1980s income inequality in the UK was higher in terms of tax units than per capita household income (Atkinson and Micklewright, 1992).

²¹Similarly, in India only approximately 3.5% of tax units filed a return in the late 1990s (Banerjee

I and II differ in how I account for this non-filing. In the family-size adjustment (Part I), I remain as close as possible to Piketty and Saez (2007b) who use census information to account for non-filing and thus define top fractiles relative to the entire US population of *potential* taxpayers (as opposed to those tax units who happen to file a return).²² The analysis of the association between capital and labour incomes (Part II) ignores the issue of non-filing, because I would have to make an arbitrary assumption about the income composition of the non-filers. Furthermore, such an adjustment might not affect the results substantially because the fraction of filers is quite high over the time I am analysing (between 88% and 96%, see Table 2.1), compared with 10-15% before WWII (Piketty and Saez, 2007b).

2.2.2 Income concept

Income is defined as (taxable) gross income excluding social security and unemployment insurance benefits. The precise definition differs between Parts I and II. Part I replicates the income definition of Piketty and Saez (2007b) and Strudler et al. (2005). In Part II, I use the sum of all component incomes reported on the tax return. The difference between the two total income variables is small (never more than 2% of income for various top fractiles, Table 2.1 panel E) and could be positive or negative. Piketty and Saez (2007b) note that this income could be alimony (paid and received) or “other income”, which refers to income from miscellaneous sources.²³ They exclude this residual income from their income composition analysis, but include it when defining fractiles. I could have (arbitrarily) allocated this income to labour or non-labour income, but this would misstate the factor income shares.

In Part I, income is defined as market income excluding all transfers. Specifically, it is calculated as AGI minus taxable unemployment and social insurance benefits, plus

and Piketty, 2010).

²²The Piketty and Saez (2007b) approach differs from Feenberg and Poterba (2000) who also base their analysis on tax units, but use the adult population growth rate in order to adjust for non-filing. This is problematic given that the number of adults per tax unit has declined over time.

²³In addition, they mention that it might refer to taxable social security and unemployment benefits. However, in the Piketty and Saez (2007b) (and my) definition of total income these have been dropped, so this cannot explain the differences. It appears that this part of the data description was copied from Piketty and Saez (2001) without accounting for the change in the income definition.

any adjustments to AGI. AGI is defined as (gross) income reported on the tax return minus some (“above-the-line”) adjustments.²⁴ These adjustments change depending on current legislation and I use Emmanuel Saez’ estimation code (with small corrections) to add them back. Social security and unemployment insurance benefits are excluded because (1) the rules for reporting them on tax returns have changed over time, and (2) they represent transfer rather than market income.²⁵ My approach differs from other analyses which use AGI or taxable income.²⁶

There exists a slight inconsistency in how payroll taxes (contributions to social security and Medicare) are treated between employees and self-employed. Both employers and employees pay an equal share of payroll taxes. For an employee, the gross income reported on the tax return is after employer-contributions but before employee-contributions. The self-employed are liable for both portions of the payroll tax, but the employer-portion is tax-deductible, as an “above-the-line deduction”. Piketty and Saez (2007b) add these deductions back into income, and hence the income variable used for the self-employed is before *any* payroll taxes. I follow this approach because I try to be consistent with Piketty and Saez (2007b). Furthermore, in Part II it is unclear how to split the self-employment payroll tax between capital and labour.

I include realised capital gains throughout, adjusted for changes in legislation affecting the taxable portion of capital gains. From 1942 to 1978, only 50% of realised capital gains were incorporated in AGI; from 1979 to 1986, 40% of capital gains were included; after 1987 all capital gains were taxable. To adjust for this, the capital gains included in AGI are multiplied by a factor of 2 between 1960 and 1978, and by a factor of 2.5 between 1979 and 1986. These exclusions only concerned long-term cap-

²⁴These adjustments include moving expenses, dividend exclusion (before 1986 up to \$200 per year in dividends could be excluded from tax), IRA contributions, self-employment tax etc. The adjustments become less important for higher incomes (Piketty and Saez, 2007b).

²⁵Social security and unemployment insurance benefits were not taxable until 1984 and 1987 respectively. Hence they were not reported on tax returns. Furthermore, currently not all social security benefits are taxable, so tax returns offer incomplete information.

²⁶Taxable income is defined as AGI minus personal exemptions and the standard or itemized deductions (some expenses incurred in earning income). Feenberg and Poterba (2000) and Slemrod (1996) use AGI and adjust for changes in tax legislation over time. My definition seems more consistent, since it excludes all transfers, rather than AGI which includes taxable social security benefits but excludes other transfers.

ital gains, which account for the majority.²⁷ Throughout this chapter, I include capital gains both in computing shares and ranking. Piketty and Saez (2007b) present top shares which include capital gains but exclude capital gains when ranking tax units. They argue that capital gains are volatile, as the result of stock price fluctuations and lumpy realisations. They point out that including gains in ranking would lead to top income shares which are no smaller, which is true empirically (Figure 2.1) and can also be shown formally.²⁸ For simplicity, I include capital gains in both ranking and computing shares.

2.2.3 Income shares as a measure of inequality and control totals

The measure of inequality used in this chapter is the share of income accruing to a particular top fractile of the income distribution, such as the top 1%. Specifically, the top income share is obtained by dividing the total income of a particular top group by total income reported in the data, which is called the control total and which may include an adjustment for non-filers as explained below. Of course, the advantage of these measures is that they focus on changes at the very top of the distribution, which have been missed by other estimates of inequality, such as the P90/P10 ratio (Bell and

²⁷The definition of long-term also changed over time. Short-term capital gains do not receive this favourable tax treatment. Piketty and Saez (2007b) argue that the majority of capital gains are long-term.

²⁸I define $S_i = \frac{\sum_{j=1}^N y_j \times \mathbb{1}\{y_j \geq y_i\}}{Y}$ as the share of total income including capital gains of top fractile i . Similarly the share of income including capital gains, but ranked according to income excluding gains, is defined as $\tilde{S}_i = \frac{\sum_{j=1}^N y_j \times \mathbb{1}\{\tilde{y}_j \geq \tilde{y}_i\}}{\tilde{Y}}$, where $\tilde{y}_j$ is income excluding capital gains. Because the incomes of those observations that are in the top fractile whether or not capital gains are included cancel out, the following relationship holds

$$S_i - \tilde{S}_i = \frac{\sum_{j=1}^N y_j \times \mathbb{1}\{y_j \geq y_i \cap \tilde{y}_j < \tilde{y}_i\}}{Y} - \frac{\sum_{j=1}^N y_j \times \mathbb{1}\{y_j < y_i \cap \tilde{y}_j \geq \tilde{y}_i\}}{\tilde{Y}} \geq 0$$

This is weakly positive because in the first term all y_j 's are at least as big as the cut-off level y_i , while in the second term they are strictly less than y_i . Furthermore, the number of tax units summed in the two terms are equal, because the number of tax units in the top 1% is the same whether including or excluding capital gains. This has the intuitive explanation that the incomes including capital gains of the “entrants” (i.e. in top fractile of income including gains, but not in top fractile of income excluding gains) are greater than the incomes including capital gains of the “dropouts” (i.e. in top fractile of income excluding gains, but not in top fractile of income including gains). The expression is zero if the two rankings coincide, i.e. $S_i = \tilde{S}_i$. The notation follows Section 2.4.

Van Reenen, 2013). Leigh (2007) notes that top shares also obey three basic principles of inequality measurement: First, they are scale-invariant, i.e. they are unaffected by an equal proportional income change affecting all incomes. Second, they obey the principle of population, i.e. merging two identical distributions would leave income shares unchanged. Third, they are anonymous, i.e. characteristics other than income leave them unchanged. However, they only weakly satisfy the Pigou-Dalton transfer principle: A top share never increases as the result of a transfer from a rich to a poor person but it stays the same if both persons are within or outside the top fractile.

Because of non-filing, total income reported on tax returns understates true total income. As explained above, the number of non-filers is estimated using census data. The literature has used two methods to adjust total income: First, one could use external information, such as national accounts income, and compute the control total as a fixed fraction of national accounts income. The main problem with this approach is that income is defined differently in national accounts and tax records, such as the treatment of imputed rent, transfers, and capital and inventory adjustments. Second, I could assign each non-filer a fraction of average reported income. Piketty and Saez (2007b) use the second approach after 1944 and the first method before that, when the fraction of filers was substantially lower. In Part I, I follow them and allocate each non-filer 20% of average reported income, which leads to a very similar control total.²⁹ I choose the second approach, because (1) the fraction of filers is relatively high during the period under investigation, (2) it does not involve the use of national accounts with different income definitions, and (3) as Piketty and Saez (2007b) argue there exist some inexplicable short-run fluctuations in the ratio of the control total to national accounts income. In Part II, I simply use total income reported in the tax file, because I do not adjust for non-filers.

²⁹Excluding capital gains, my control total is on average 0.029% lower than the control total in Saez (2013). Across the 43 years, the percentage difference ranges from -1.232% to 1.065%. Including capital gains, it is on average 0.058% higher than Saez (2013) ranging from -1.112% to 1.501% across the years. The slightly bigger difference in the control totals when including capital gains can be explained by the fact that Piketty and Saez (2007b) assume that non-filers do not have any capital gains, whereas I assume that they realise 20% of average capital gains.

2.2.4 Summary

In summary, the IRS data are appropriate for my analysis because they (1) oversample at the top and thus provide a sufficient sample size; (2) have effectively universal coverage at the top, whereas high income households generally have lower response rates in surveys (Feenberg and Poterba, 2000);³⁰ (3) do not suffer from top-coding, in contrast to the CPS; (4) are not subject to the same recall errors as surveys, because they are derived from administrative records; (5) have a comprehensive measure of market income, including realised capital gains and stock options. The results need to be interpreted carefully because the data (1) relate to tax units, not households as in the case of most surveys; (2) are subject to short-run income shifting in response to changes in the tax code, which is one reason for focussing on long-run changes (see Appendix); (3) include only taxpayers, which makes them inappropriate for analysing the bottom of the distribution; (4) include only taxable income, which ignores income sources particularly important at the bottom of the distribution.

As I already mentioned, the data construction is slightly different between Parts I and II. Part I of the chapter follows the approach by Piketty and Saez (2007b) as closely as possible, adjusting for tax unit size. I adjust for non-filers using the total number of potential tax units from Saez (2013) and population estimates from the US Census Bureau. I assign them 20% of average income to adjust the control total. I conduct the analysis both at the tax unit and the individual level.

In Part II (the analysis of the association between capital and labour), I do not adjust for non-filers, and income is defined as the sum of component incomes, excluding “other income”. Capital gains are included throughout. I drop observations which are negative in any of the factor income variables labour, capital (including capital gains) and business income.³¹ Although this assumption is necessary in my analysis, it is important to stress that it may not be uncontroversial. For instance, high positive business or capital incomes are often justified by the downside risks involved, so in

³⁰This issue is discussed in detail also in the chapter on global inequality.

³¹In the analysis, business income is split two-thirds to labour and one-third to capital. Negative observations are dropped before this because otherwise the labour share would be reduced artificially.

principle one should account for the downside as well as the upside. It would be easier to deal with the negative incomes if it were possible to link the data over time. In other words, while it is not surprising that some tax units will have negative business or capital incomes in some years, these are unlikely to be the same tax units all the time. Negative incomes are “problematic” for the factor income decomposition, because the top shares of the distributions of factor incomes could be greater than 1.³² Dropping negative incomes follows Auerbach and Hassett (2002) who also drop observations with negative AGI. Alternatively, I could have set the negative incomes to zero but this would involve changing the data. Instead, I prefer to restrict the analysis to the sub-sample with positive incomes. Furthermore, the trend in the top shares of the distributions of capital and labour incomes reported in Figures 2.5, 2.6 and 2.7 (bottom left panel) is very similar when negatives are recoded to zero. Note that the issue of tax units with negative income is not the same as non-filing, because otherwise a negative income would be replaced by the income imputed for non-filers, which is positive.

2.2.5 Summary statistics

Table 2.1 presents the summary statistics for the main sample across a number of years (focussing on the Tax Reform Act of 1986 (TRA86) which is described on p. 89), taking sampling weights into account. The number of observations in the main sample is around 100,000 across the years. The fraction of tax units filing a tax return is high at around 95%, so I do not expect that ignoring non-filers in Part II would affect the results in a major way. In panels A and B, I report average and threshold incomes excluding and including capital gains for a number of top fractiles deflated to 2005 USD.³³ Non-filing is accounted for in the definition of fractiles (the non-filers are assumed to be at the bottom of the income distribution as explained in Section 2.3.1)

³²For example, when keeping the negative incomes, the share of capital income (including one third of business income, as explained on p. 59) that accrues to the top 10% of capital income recipients exceeds 100% in the early 2000s.

³³The CPI data are taken from Saez (2013). Until 1978, he uses the CPI index in the Economic Report of the President. After 1978, the CPI-U-RS is used, which is the version of the CPI computed under a constant methodology, and which is therefore more comparable over time. The CPI-U-RS rises slower than the standard CPI and thus shows a higher growth in real incomes (BLS, 2013).

and the overall average income (by imputing 20% of average income). For the years in which the size adjustment is possible, I also report the average tax unit size, the fraction of tax units who are dependents and their average income (panel C).

In Part II, I drop observations whose labour, capital (including capital gains), or business income is negative. Panel D of Table 2.1 reports the fraction of negative observations in either of these income sources. Not surprisingly, labour incomes are weakly positive, so the restriction on labour incomes is inconsequential. In most years, some tax units have both negative business and capital incomes. The sample restriction affects around 10% of the (weighted) sample (see the row “fraction dropped”), whose average income is greater than the average, but less than the top 10% threshold. The bottom rows of panel D report the number of observations used in Part II.

Finally, in panel E, I summarise “other income” which is not included in labour or capital, and thus excluded in Part II. This income accounts for less than 1% of income overall, and around 1% amongst the top fractiles (see also Piketty and Saez, 2007b).

Part I Top income shares adjusted for tax unit size

The first step in this analysis involves reconstructing the top income shares according to the Piketty and Saez (2007b) specification, which I refer to as the baseline estimates. Table 2.B.1 (without capital gains) and Table 2.B.2 (with capital gains) compare my baseline shares with the updated shares from Saez (2013), which between 1966 and 2001 are derived from the same data source.³⁴ They are very similar, but there are some differences. As mentioned above, Piketty and Saez (2007b) assume that non-filers do not have any capital gains, whereas I assign them 20% of average capital gains. Furthermore, they rank observations by AGI, while gross market income is used to compute the income share. I use gross market income both in ranking and in computing shares.³⁵

2.3 Adjusting for tax unit size

Throughout this chapter, the unit of analysis is the tax unit, which could consist of a number of taxpayers and dependents. Ultimately one is concerned with the welfare of the individual, not the tax unit. Nevertheless, a larger reference unit might be appropriate because of income-sharing. Otherwise, many individuals such as children

³⁴Although micro data were available, Piketty and Saez (2007b) use the tabulations in 1960, 1962 and 1964 because the early micro-files only contain few variables. For updating the series, Saez (2013) uses tabulated data after 2001.

³⁵Ranking by AGI (not shown here) further reduces the average differences (in absolute terms) between my top shares (excl. capital gains) and Saez (2013), with the exception of the two very top fractiles. The average difference for the P99.9-99.99 (P99.99-100) is -0.07pp (-0.17pp) compared with 0.00pp (for both) in Table 2.B.1.

or spouses working at home would have zero incomes. The implicit assumption is that people share income within the tax unit, which is appropriate because this is a legal requirement for claiming dependents. As mentioned on p. 37, the proportion of married tax units declined slower at the top than the rest of the population. This would have increased top income shares by itself because single, and thus poorer, tax units are added faster at the bottom than the top.

2.3.1 Empirical specification

The IRS data report the number of personal exemptions, which are fixed amounts that can be deducted from taxable income for each taxpayer and dependent. Tax unit size is defined as the number of dependents plus the primary or secondary taxpayers. For confidentiality reasons, the total number of dependents is limited to three for high-income tax units in the most recent rounds of the data.³⁶ I implement this restriction in all years and hence my tax unit size variable is a lower bound on the true size. The tax units which are dependents of another tax unit are assigned a size of one (there is still one person!).³⁷ Between 1984 and 1992 the primary and secondary taxpayer exemptions for top-income tax units are not reported in the data because of data protection procedures. Hence I exclude these years from the analysis which is thus from 1960 to 1983 and from 1993 to 2005.³⁸ Hungerford (2013) also measures tax unit size by the number of personal exemptions but he does not account for top-coding of dependent exemptions for high-income tax units, or the suppression of the taxpayer exemptions before 1993.³⁹ In their investigation of the marriage penalty, Auerbach and Hassett (2002) adjust for the number of adults (i.e. whether single or joint filers), using both per capita incomes and a square-root equivalence scale.

³⁶Since at least 2000, the number of children living at home is limited to a maximum of three for high income earners (\$200,000 AGI or more).

³⁷These tax units do not claim any exemptions. By default they would therefore have a size of zero. This would make size-adjusted income missing, since I divide by zero. See p. 43 for a detailed discussion of how to treat these observations.

³⁸I could have imputed the number of taxpayers from the marital status, but the IRS also recoded marital status for top incomes.

³⁹Furthermore, his analysis is quite different, as he concentrates on the Gini index, its factor income decomposition and begins in 1991.

One can compare tax unit size with household size according to the Census definition (see Section 2.2.1 for a comparison of tax units and households). On average, there are fewer people in a tax unit than a household (and a family is bigger still).⁴⁰ Between 1970 and 2005, the average household size reported by the US Census Bureau is 2.76, which compares with 2.19 for my data. Over the years, the difference lies between 22.2% and 29.9%. These differences can be explained as follows: First, as discussed in Section 2.2.1, there are definitional differences between tax units and households. Second, the Census refers to the number of own children whereas my size measure counts the number of dependents, which might include dependent parents. Furthermore, children are defined differently in the Census and the IRS, which would also count dependent grandchildren or nephews (Mann and Salvo, 1984). Third, the IRS estimate of dependents is top-coded, so I would expect it to be an underestimate.

Consumption needs increase with tax unit size, but due to scale economies in consumption and production, in a less than proportionate way. I use a square root equivalence scale (defined below) to adjust household incomes. Many different equivalence scales have been used in the literature, which imply different assumptions about how household needs increase with its size (Buhmann et al., 1988).⁴¹ I use the square root scale, because in the data I cannot differentiate between types of dependents as required by scales which weigh adults and children differently. This particular equivalence scale has been called an “intermediate scale” (Atkinson et al., 1994). It has been used in recent OECD publications (OECD, 2008) and in the USA with the Consumer Expenditure Survey (Fisher et al., 2013), the IRS data (Hungerford, 2013), the CPS (Burkhauser et al., 2009), and a combination of both the IRS and CPS data (CBO, 2012).

I have conducted the analysis both at the tax unit (as before) and the individual

⁴⁰Families are related by birth, marriage or adoption (footnote 19). Between 1970 and 2005, the average family size is 3.28.

⁴¹Atkinson et al. (1995) use a square root scale but compare their results with the no adjustment case and a per capita adjustment. They find that in general the square root scale returns the lowest levels of inequality. Alternative equivalence scales change the inequality rankings of countries depending on the joint distribution of household size and income. The broad cross-country patterns, however, remain unchanged.

level. In the first case, total income in the data is “too low” because I divide tax unit income by the square-root of tax unit size. Conversely, total income is “too high” in the second case, where tax unit income is divided by the square-root of tax unit size as part of the equivalisation, and then multiplied by size in order to translate the analysis to the individual-level (i.e. overall I multiply tax unit income by the square-root of tax unit size). One of the reason for adjusting for tax unit size is that I am concerned with individual welfare, so the second option seems more appropriate.

In the analysis at the tax unit level, non-filers are defined relative to tax units using the number of *potential* tax units by Saez (2013). Potential tax units are estimated from census data and the CPS as the sum of married men, single men and women aged 20 and over, and divorced and widowed men and women. A tax unit thus consists of a married couple living together or a single adult (plus any dependents in both cases). This approach assumes that all married couples file a joint return. This is a reasonable assumption, since in most cases joint filing reduces the overall tax liability for a married couple and hence the proportion of married couples filing separately has been small in practice.⁴² The share of married couples filing separately is on average 2.08% of tax records. It has declined over time and it is smaller in the top fractiles compared with all filers.⁴³ The fraction of filers is estimated as the ratio of the total returns in the data (i.e. the sum of the sampling weights) to the total potential returns estimated from census data. It is assumed that everyone who does not file a tax return is at the bottom of the

⁴²For a married couple, filing jointly is typically advantageous, because there are more tax credits and benefits and it attracts lower tax rates. For example, in a couple where one spouse has no or a much lower income, tax liabilities are reduced because the combined income rather than the individual incomes are taxed. In some cases, however, separate filing could be better. For instance, in order to claim unreimbursed medical expenses or miscellaneous itemised deductions, these expenses need to cross a threshold level defined as a percentage of taxable income. Hence reducing taxable income by filing separately would reduce this threshold. Furthermore, a married couple might file separately, if it is getting divorced or if one spouse suspects the other of tax fraud. Note that there also exists the issue of a “marriage tax” or “marriage penalty”. In some cases total taxes can increase when a couple gets married, because the tax brackets for single and married people are different. However, this compares single and married taxes, not the taxes for a married couple filing jointly or separately. The “married, filing separately” brackets are exactly half of the “married, filing jointly” schedule.

⁴³The share declined from 3.77% in 1968 to 0.88% in 1985. Since 1990 it was approximately 2% (less than 2% since 1999). The share in the four top shares (top 0.5%, 1%, 5% and 10%) is everywhere less than the share amongst all records. Since 1980, the share in the top 0.5% and the top 1% is greater (average of 1.3% and 1.1%) than the share in the top 5% and top 10% (average of 0.72% and 0.67%), while before they were more similar.

distribution.⁴⁴ Hence, in order to define the top income groups (e.g. the top 10%), the adjusted cumulative frequency $\widetilde{F}(w)$ is used. It is computed as $\widetilde{F}(w) = 1 - \theta[1 - F(w)]$, where θ is the fraction of filers and $F(w)$ is the raw cumulative frequency computed amongst the tax returns filed (using sampling weights). In the construction of the control total, I assign each non-filing tax unit 20% of average individual equivalent income, which follows the baseline shares estimated according to Piketty and Saez (2007b). Throughout I assume that there are no systematic differences in tax unit size between filers and non-filers.⁴⁵

In the second case, each tax unit is weighted by its size in addition to the sampling weight, i.e. the total weight is the product of the sampling weight and tax unit size. The fraction of filers is now defined relative to the US population, obtained from the Population Division of the US Census Bureau.⁴⁶ The fraction of filers is very similar in the two cases.⁴⁷ Given the top coding of dependents, I would have expected the fraction of filers to be lower in the second case. This is counteracted by tax units with no exemptions who are essentially double counted. As before, the control total is derived by allocating 20% of the average individual equivalent income to each non-filer.

2.3.2 Results

Figures 2.3 and 2.4 compare the size-adjusted top income shares with those reproduced following the Piketty and Saez (2007b) methodology. In the literature review, I already discussed the strong increase in top income shares between the 1960s and the 2000s. The abrupt jumps in the mid-1980s are likely to be related to changes in the tax code as part of TRA86 (discussed in detail in Section 2.A.3). For the richest tax units, TRA86

⁴⁴To be precise, the analysis assumes that the non-filers are outside the top fractile.

⁴⁵In order to claim benefits through the Earned Income Tax Credit, tax units have to file a return. Therefore, it might be argued that low-income non-filing tax units have fewer dependents. However, it is not clear what the appropriate assumption about their dependents should be. Furthermore, participation in the EITC by eligible tax units is far from complete, although by definition it is difficult to obtain data on the non-participants (Holt, 2006).

⁴⁶The figures for 2000-2005 are obtained from the “Annual Estimates of the Resident Population for the United States”, released December 2009. The population numbers for 1960-1999 are from the “Historical National Population Estimates: July 1, 1900 to July 1, 1999”, revised June 28, 2000.

⁴⁷In terms of tax units, the fraction of filers ranges from 87.99% to 96.25%, with an average of 93.49%. For the individual-level analysis, the fraction of filers is slightly lower with an average of 91.63%, ranging from 84.82% to 94.30%.

increased taxes on capital gains and brought the personal income tax rate below the corporate rate. A tax-reducing shift of rich people's incomes from corporations to persons probably contributed to the strong increase in the mid-1980s of the top 1% share excluding capital gains (Figure 2.3, top left panel). Anticipation of lower tax rates may also have caused receipts of income, such as the exercise of stock options, to be delayed from 1986 to 1987, making the top 1% share in 1986 artificially low. By contrast, the increase in capital gains taxation provided incentives to realise these gains earlier, which could explain why the jumps in the top 1% share are less abrupt when capital gains are included (Figure 2.3, bottom left panel). Because this type of fiscal manipulation can have important effects in the short-run, I focus on the long-run trends which are more likely due to real changes (Piketty and Saez, 2007b).

The results on the size-adjusted top income shares have to be interpreted carefully because the different fractiles do not refer to the same people due to re-ranking induced by the equivalisation. In 2000, 0.2% of tax units are in the top 1% according to baseline incomes, but not after adjusting incomes for tax unit size.⁴⁸ Accounting for tax unit size reduces inequality but leaves the trend unaffected as the three series track each other very closely. This confirms evidence from household surveys where the inequality of equivalent incomes is often smaller than without an adjustment for household size (Atkinson et al., 1995). The equivalisation affects the top 10% share more than the top 0.5% share.⁴⁹ This suggests that it affects overall inequality more than top inequality. This might be because tax unit size is more homogeneous at the top of the distribution. Before 1995, the inequality among individuals is smaller than the inequality among tax units. Again, this is similar to evidence from household surveys (Atkinson et al., 1995). After 1995, this ranking reverses with the exception of the top 10%, although the differences are quite small. It is not clear why this reversal happens.⁵⁰

⁴⁸In other words, roughly 20% of tax units who are in the top 1% of baseline incomes drop out of the top 1% after the equivalisation. This conditional probability is slightly lower when including capital gains. It has fallen over time from a peak of 25% in the 1970s. It increases for larger fractiles such as the top 10% (to around 23%), and it falls for the smaller fractiles such as the top 0.01% (to around 16%).

⁴⁹In absolute terms this can easily be seen from the chart. This is also true proportionately for the equivalisation at the tax unit level (except for around 2000) and the individual level after the 1980s.

⁵⁰The dependent tax units might have contributed, but there is no conclusive evidence. There were

In summary, the observed changes in the top income shares are robust to controlling for tax unit size.⁵¹ Furthermore, the adjustment does not make much difference to the level of the top shares. Hence, I follow the approach in the literature and omit the size adjustment in the rest of the chapter.

no dependent tax units before 1984 (with the exception of very small proportions in 1962 and 1964). The proportion of dependent tax units peaked between 1997 and 2000 (at around 9%). The share of the top 0.5% who are dependents peaked in 2000 (at around 0.26%), but it is not substantially higher than in 1994. In 2002, the proportion of dependents amongst the top 5% is 0.05% which is lower than in 1993, 1994, or 1995.

⁵¹In Canada, the increase in top shares is very similar whether measured on an individual or tax unit level (Saez and Veall, 2007). In the US, Kopczuk et al. (2010) show that the individual-based top *wage* shares derived from Social Security records evolved very similar to the tax unit analysis.

Part II Association between labour and capital income

This part of the chapter analyses how the association between labour and capital income has changed over the period 1960-2005. As shown by Figure 2.2, the labour share at the top of the income distribution increased over this period. I am interested in whether capitalists and rentiers have remained separate classes, or whether individuals at the top of the income distribution increasingly receive income from both sources, i.e. the association of factor incomes increased.

As discussed in Section 2.2 (summary in Section 2.2.4), the data construction is different from Part I of the chapter: Total income is defined as the sum of component incomes; there is no adjustment for non-filers, i.e. fractiles are defined relative to the population of filers; capital gains are included throughout; observations with negative labour, capital or entrepreneurial income are dropped. The adjustment for tax unit size left the evolution of top income shares largely unchanged, so I omit it in this part.

Table 2.B.3 shows the effect on top shares of using the sum of component incomes and not accounting for non-filing.⁵² All shares are within 1pp of each other. Across the five fractiles, the average difference ranges from 0.33pp to 0.05pp. One can see the effect of changing the sample (while keeping the total income definition and the non-filing adjustment) by comparing Tables 2.B.2 (column “own”) and 2.B.3 (column “old”). These differences between the top shares are generally small, although they are bigger

⁵²Table 2.B.3 uses the Part II sample throughout, which excludes observations with negative labour, capital or business incomes. For each top share, the first column adjusts for non-filers and uses the standard total income definition of gross market income. The second column follows the methodology in Part II, i.e. it computes top shares without an adjustment for non-filing and uses the sum of component incomes as total income.

at the very top (for the top 0.1%, the difference is 0.83pp on average with a maximum of 2.18pp in 1986). Excluding negative observations (the change in the sample) has a greater effect than excluding the other income category and not accounting for non-filers. However, the different series track each other very well, so there seems to be no evidence of any systematic bias resulting from the change in the data construction.

I allocate income to two factors, labour and capital. Given the difficulties in thinking about associations in two or more dimensions, this offers a reasonable starting point. Furthermore, association matrices can only be presented in two dimensions, while the factor income decomposition and a full parametric copula estimation could obviously handle more. Business income (Schedule-C and partnership income) has been allocated two thirds to labour and one third to capital. The weights used in splitting business income are essentially arbitrary. My approach is almost the same as Johnson (1954), who attributes 65% of non-farm entrepreneurial income to labour. In a more careful approach, for which my dataset lacks the necessary information, Feinstein (1968) assigns every self-employed individual the average wage in the same sector of the economy, and the residual income is defined as the return to capital.

There exists a related literature in macroeconomics, which is concerned with measuring labour and capital shares in national accounts. The existence of intermediate institutions and differences in the income definitions make it difficult to move between personal and national income (Atkinson, 2009). In the US, the share of total employee compensation in national income was fairly stable in the second half of the 20th century, at on average 62% in the 1960s, 66% in 1970s and 1980s, 65% in 1990s (Feldstein, 2008), but declined to 58% in recent years (Elsby et al., 2013). However, this measure of the labour share excludes any self-employment income (Elsby et al., 2013). In many cross-country analyses, labour shares of developing countries, where self-employment is more common, are systematically understated, because self-employment income is allocated to capital income and labour is approximated by employee compensation, as argued by Gollin (2002). His revised labour share (when self-employment income is split between labour and capital) lies between 0.654 and 0.686 on average which is

similar to my weight of two-thirds.

In summary, labour income is defined as the sum of wages, pensions, and two-thirds of business income (which in turn is defined as the sum of Schedule-C and partnership income). Capital income is defined as the sum of dividends, (taxable) interest, rents, estate income, royalties, capital gains (adjusted for the changes in legislation), S-corporation income, and one-third of business income. As explained on p. 94, allocating S-corporation income to capital instead of business income, is different from Piketty and Saez (2007b) (but similar to CBO, 2012), and is justified in order to minimise the effects of income shifting as the result of TRA86.

2.4 Factor income decomposition: The alignment coefficient

I first present a factor income decomposition of top income shares. This is a formal derivation and (the first) empirical application of a decomposition suggested by Atkinson (2007b), who builds on Meade (1964).⁵³ It is closely related to factor income decompositions of other inequality measures, such as the Gini index (Shorrocks, 1982; Lerman and Yitzhaki, 1985). Like these decompositions, inequality is decomposed into three elements: The share of each factor in total income, the inequality of the distribution of income from each of the factors and the association between incomes from different factors.

The top share can be written as $S_i = \frac{Y_i}{Y}$, where Y is total income in the data. The term Y_i denotes total income of tax units with income greater or equal to y_i , i.e. $Y_i = \sum_{j=1}^N y_j \times \mathbb{1}\{y_j \geq y_i\}$, where N is the total number of observations (ignoring any non-filing). In other words, if one is interested in the share of the top 1%, the cut-off value y_i would be the income of the 99th percentile. Total income is derived from M factors, i.e. for any individual j , $y_j = \sum_{m=1}^M x_{j,m}$. Total income, in say the top 1%, is the sum of the incomes from each of the factors. Therefore, the top share can be

⁵³Atkinson (2007b, p. 35) states Equation 2.4 in words without any derivation or empirical application.

written as

$$S_i = \frac{Y_i}{Y} = \frac{\sum_{m=1}^M \widetilde{X_{i,m}}}{Y} = \sum_{m=1}^M \frac{\widetilde{X_{i,m}}}{Y} \quad (2.1)$$

where $\widetilde{X_{i,m}} = \sum_{j=1}^N x_{j,m} \times \mathbb{1}\{y_j \geq y_i\}$. The term $\widetilde{X_{i,m}}$ denotes, for example, total labour income received by the top 1% of income recipients. It is important to note that the top fractile is still defined in terms of total income, not factor income. Expanding the fractions by total income from factor m , $X_m = \sum_{j=1}^N x_{j,m}$, I can further rewrite the top share as

$$S_i = \sum_{m=1}^M \frac{\widetilde{X_{i,m}}}{Y} = \sum_{m=1}^M \frac{X_m}{Y} \times \frac{\widetilde{X_{i,m}}}{X_m} \quad (2.2)$$

The first term is the factor share, i.e. the proportion of total income which is derived from income source m . It can also be written as the ratio of the mean of income from factor m to the mean of total income $\frac{\mu_m}{\mu}$. I will refer to the last term as the pseudo-share ($\widetilde{S}_i^m = \frac{\widetilde{X_{i,m}}}{X_m}$) as in Shorrocks (1982). It denotes the share of income from factor m received by the top 1% of total income recipients. The decomposition in Equation 2.2 is the expression given in Meade (1964) [p. 28]. Expanding the fraction by $\frac{X_{i,m}}{X_m}$, gives the expression in Atkinson (2007b)

$$S_i = \sum_{m=1}^M \frac{X_m}{Y} \times \frac{\widetilde{X_{i,m}}}{X_m} = \sum_{m=1}^M \frac{\mu_m}{\mu} \times \frac{X_{i,m}}{X_m} \times \frac{\widetilde{X_{i,m}}}{\frac{X_{i,m}}{X_m}} = \sum_{m=1}^M \frac{\mu_m}{\mu} \times S_i^m \times \frac{\widetilde{S}_i^m}{S_i^m} = \sum_{m=1}^M \frac{\mu_m}{\mu} \times S_i^m \times A_i^m \quad (2.3)$$

where $X_{i,m} = \sum_{j=1}^N x_{j,m} \times \mathbb{1}\{x_{j,m} \geq x_{i,m}\}$. The term $S_i^m = \frac{X_{i,m}}{X_m}$ is the share of income from factor m that accrues to the top fractile i of recipients of income from factor m . For example, this is the share of labour income going to the top 1% of the labour income distribution. It describes the inequality in the marginal distribution, as measured by the top share. Note that $S_i^m \geq \widetilde{S}_i^m$.⁵⁴ For example, this is because

⁵⁴The reasoning is as in footnote 28. Because the incomes of tax units who are in the top fractile i of both income from factor m and total income ($= \frac{\sum_{j=1}^N x_{j,m} \times \mathbb{1}\{x_{j,m} \geq x_{i,m} \cap y_j \geq y_i\}}{X_m}$) cancel out, I can write

the labour income of those tax units who are in the top fractile of labour but not total income is greater than the labour income of those tax units who are in the top according to total but not labour income. The last term (A_i^m) denotes “the share in earnings (capital income) of the top 1 percent of income recipients divided by the share of top 1 percent of earners (capital income recipients)” (Atkinson et al. (2011), p.61). Atkinson (2007b) calls this the “alignment coefficient”. Considering only two sources of income, labour and capital, I can write the decomposition as

$$S_i = \frac{\mu_L}{\mu} \times S_i^L \times A_i^L + \frac{\mu_K}{\mu} \times S_i^K \times A_i^K \quad (2.4)$$

where μ_L is mean labour income and μ is mean total income, so $\frac{\mu_L}{\mu}$ is the share of total income from labour. S_i^L is the share of labour income accruing to the top fractile i of the labour distribution, i.e. it is the share of labour income accruing to say the richest 1% of labour income recipients. As explained in more detail below, A_i^L is the alignment coefficient between labour income and total income for top fractile i . The capital terms are defined similarly.

The alignment coefficient captures the connectedness between the distribution of income from factor m and the distribution of total income. It lies between 0 and 1, the upper bound following from $S_i^m \geq \widetilde{S}_i^m$ (see footnote 54). If the top income recipients receive no labour income, $\widetilde{X}_{i,L} = 0$, and hence $A_i^L = 0$. On the other hand, if everybody in the top fractile of the (total) income distribution is also to be found in the top fractile of the distribution of labour income, then $\widetilde{S}_i^L = S_i^L$ and thus $A_i^L = 1$. A labour alignment coefficient of 0.9 means that 90% of total earnings of those in the top 1% of earners go to tax units who are also in the top 1% of total income. Hence only 10% of these earnings go to tax units in the bottom 99% overall. In other words, if the top 1% of earners receive 11% of earnings and the alignment coefficient is 0.9, then as much as 10% of that 11% goes to tax units who are in the overall top 1%.

$$S_i^m - \widetilde{S}_i^m = \frac{\sum_{j=1}^N x_{j,m} \times \mathbb{1}\{x_{j,m} \geq x_{i,m} \cap y_j < y_i\}}{X_m} - \frac{\sum_{j=1}^N x_{j,m} \times \mathbb{1}\{x_{j,m} < x_{i,m} \cap y_j \geq y_i\}}{X_m} \geq 0$$

These figures are approximately equal to the values observed between 1990 and 2005 for the top 1% (Figure 2.5).

The alignment coefficient does not differentiate within the bottom 99%, i.e. it does not tell us how much of the earnings of the top 1% of earners goes to people in the next 1% (of total incomes). Furthermore, the alignment coefficient depends on changes in the marginal distributions. Moreover, changes in the distribution of income from capital affect A_i^L through the pseudo-share in the numerator which includes total income. Dardanoni and Lambert (2001) show in their paper on horizontal equity in taxation that only a rank-based measure is invariant to all monotone transformations in the marginals. As I will show in Section 2.5, a copula-based measure of association satisfies these conditions.

2.4.1 Results

Figure 2.5 shows the results of the factor income decomposition for the top 1%: The top left panel shows the share of total income going to the top 1% of total income recipients, i.e. the left-hand-side of Equation 2.4. The remaining three panels refer to the different components of the decomposition. Of course, the factor shares add up to 1, so I only show the labour share, and are the same for all top fractiles. The bottom left panel shows the inequality in the marginal distributions as measured by the share in total labour (capital) income of the top 1% of labour (capital) income recipients. The bottom right panel is the alignment coefficient. Figures 2.6 and 2.7 show the results for the top 5% and top 0.5% respectively.

All three top shares (top left panel) peaked in 2005, after an earlier peak in 2000. In the 30 years between 1975 and 2005, the top 1% (5%) share increased by 11.75pp or 150% (by 14.26pp or 70%). The top 0.5% share increased by even more in proportional terms, by 10.27pp or 195%. The estimated level of the top share is lower than in Saez (2013) (Figure 2.1 or Table 2.B.2) but the two series track each other very closely (see p. 58).⁵⁵

⁵⁵As shown above, most of this effect comes from dropping observations with negative labour, capital or business income. These observations still have positive average income (Table 2.1, Panel

The marginal distributions (bottom left panel) have also become more unequal over this period, as measured by the share of factor income received by the top fractiles of the factor income distributions. The top capital income share in 1960 appears to be an outlier, because the tax data in that year are not very robust. The top wage share in Piketty and Saez (2007b), which is estimated differently, is slightly lower than mine, but follows a very similar trend.⁵⁶ Note that the top capital share does not drop after 2000, in contrast to the total income and labour shares. This might be due to capital gains realisations offsetting other fluctuations in capital incomes. Because capital is more unequally distributed than labour income, the labour and capital axes are quite different: In 2005, the top 1% of labour earners command 12.78% of labour incomes, whereas the top 1% of capital recipients get 63% of all capital incomes. For comparison, Atkinson (2007a) reports a share of just under 6% for the top 1% of earners in the UK in 2000. Given the large inequalities in the wealth distribution (Kopczuk and Saez, 2004), it is of course not surprising that capital incomes are more unequally distributed than labour incomes.

The absolute increase in top shares between 1975 and 2005 is greater for capital than labour. However, in relative terms, the inequality of labour incomes increased by more.⁵⁷ The relative increase in the total income top share was greater than the increase in either the labour or capital share. This suggests that the other components of the decomposition contributed towards increasing the top share. The role of the alignment coefficient is considered in more detail below.

My results show that the inequality in capital incomes increased. This appears contradictory to Piketty and Saez (2007b) who write that the “the secular decline of top capital incomes is due to a decreased concentration of capital income rather than a

D), and thus total income declines. The decline in the top share suggests that incomes in the top fractile declined even faster.

⁵⁶See Figure 5.9 of that paper. They estimate the fraction of filers from the number of tax units with positive wage income, and the control total as total compensation of employees from national accounts. I only use the information in the tax data. The Piketty and Saez approach would be difficult to implement for capital incomes because national accounts include rents, count other capital incomes as retained earnings and exclude capital gains that do not reflect undistributed profits.

⁵⁷For labour, the top 1%, top 5% and top 0.5% shares increased by 6.87pp (116%), 9.43pp (52%) and 5.67pp (151%) respectively. The corresponding figures for capital income are 28.57pp (84%), 18.87pp (28%), and 28.42pp (114%).

decline in the share of capital income in the economy as a whole” [p. 153]. However, it is important to bear in mind differences in the two analyses. First, Piketty and Saez refer to the whole period between 1916 to 1999, while my analysis starts in 1960. Second, they concentrate on the share of top incomes derived from capital (also see Figure 2.2), i.e. the pseudo-share, whereas I estimate the share of capital incomes received by the top 1% of *capital recipients*. The discussion in Bivens and Mishel (2013) also relates to a decomposition into pseudo- and factor shares, i.e. Equation 2.2. Third, Piketty and Saez show that the true top *wealth* share (i.e. the share of total wealth held by a top fractile of the wealth distribution) declined strongly from 1916 to the 1970s, and since then has increased only modestly which is similar to the results I find for capital incomes.

Ignoring the first year of data, the labour share in total income (top right panel) fluctuated by approximately 5pp between 86.9% and 81.0%. It appears to decline, although there is no clear trend given the large fluctuations. Piketty and Saez (2007b) also find that factor shares remained roughly constant over this period. On the other hand, Bivens and Mishel (2013) argue that a fall in the labour share accounts for a substantial part of the increase in the top 1% share. However, Bivens and Mishel use the CBO (2012) data, which adopts a different income definition (e.g. including transfers and in-kind benefits), and they exclude self-employment income from their labour share.⁵⁸ As a result, their labour share is substantially lower than mine.⁵⁹ The same applies to macroeconomic labour shares in national income, which are typically around $\frac{2}{3}$ (see discussion on p. 59). Using national accounts, Piketty and Saez (2007b) report an adjusted capital share in aggregate personal income of approximately 20% in the 1990s, which is quite similar to my result. On the one hand, national accounts understate the capital share because they exclude those capital gains which do not reflect undistributed profits. On the other hand, total dividends in national accounts exceed those reported on tax returns, because the former also include non-taxable

⁵⁸Elsby et al. (2013) find that the fall in the US macroeconomic labour share is explained by the exclusion of self-employment income.

⁵⁹According to Bivens and Mishel, the labour share (excluding self-employment income) of total income (including transfers and in-kind benefits) drops from 69.8% in 1979 to 60.3% in 2007.

dividends paid out through funded pension plans or retirement savings accounts.

The labour alignment coefficient (bottom right panel) for the top 1% and top 0.5% follows a very similar trend. The TRA86 appears to have led to a one-time upward shift (rather than a prolonged change) in the alignment coefficient, being relatively stable before and after that change, although it seems to have peaked in the early 1990s. For the top 5%, the rise in the labour alignment coefficient resulting from TRA86 is less abrupt. Compared with the early half of the period, the association of labour and total incomes increased for all three top fractiles. The alignment coefficient for capital fluctuates more which is not surprising given that capital incomes (including capital gains) are more volatile. For the top 5%, the association of capital has become stronger.

Note that the two alignment coefficients are measured on different scales. The labour alignment coefficient is greater than the capital counterpart for the top 5% and the top 1% in the latest years. By contrast, there is substantial overlap for the top 0.5%. As already mentioned, a value of 0.9 for the labour alignment coefficient (after 1990) implies that 90% of total earnings of those in the top 1% of earnings go to tax units who are also in the top 1% overall. For capital, the corresponding coefficient is slightly lower at around 0.86 in recent years. Hence, compared with labour income, a slightly higher share of total capital income goes to tax units who are in the bottom 99% of total incomes. In other words, top wage earners are somewhat more likely to be at the top of the total income distribution compared with top capital income recipients. Moving to the top 5% (Figure 2.6) increases the alignment coefficient for labour incomes but reduces it for capital incomes. Roughly 95% of total earnings of those in the top 5% of labour incomes are received by tax units who are also in the top 5% of total incomes. Conversely, only 76% of capital incomes in the top 5% of the capital distribution are received by those who are also in the top 5% of total income. This suggests that the tax units between P95 and P99 of the capital distribution are more likely to be outside the top 5% of total incomes compared with the top 1% of capital incomes.

The alignment coefficient is difficult to interpret because it incorporates the total income distribution. Section 2.5.1 addresses the association between capital and labour income directly as opposed to the association between the factor incomes with total income. Similar to the alignment coefficient, the association between labour and capital is found to be asymmetric: Top labour earners are more likely to also be amongst the top capital incomes, than the top capital income recipients being at the top of the labour distribution.

2.5 The copula function

As mentioned above, the association term in the factor income decomposition is not independent of changes in the marginal distributions. In this section, I introduce an analytical framework based on the copula function, which offers a clean separation of the joint distribution of labour and capital into the marginal distributions and a term capturing the association. The copula function is unaffected by a transformation in the marginals which preserves the ranks in each factor income. For example this would apply to a shift from capital to labour income which decreases capital incomes proportionately and increases labour income proportionately (Aaberge et al., 2013). Copula functions have been widely used in actuarial science to describe multidimensional risks. In economics, they have been used to study the horizontal equity of the tax system (Dardanoni and Lambert, 2001), income mobility by considering the dependence over time (Bonhomme and Robin, 2009), and the association between health and income (Quinn, 2007). Nevertheless copula functions have received relatively little attention in economics although they offer a comprehensive analytical framework to study the interactions between random variables, such as multi-dimensional inequality (Atkinson, 2011).

In my specification, total income consists of income from labour and capital. Income, X , can thus be described as a two-dimensional vector $X = (L, K)$, where L denotes wage income and K refers to capital income. The joint distribution function of X is called $H_X(l, k) : \mathbb{R}^2 \rightarrow [0, 1]$. For every x , a particular realisation of the random

vector X in $\mathbb{R}^2$, I have

$$H_X(l, k) = \Pr[L \leq l \text{ and } K \leq k] \quad (2.5)$$

The marginal distribution function of labour income, $F : \mathbb{R} \rightarrow [0, 1]$, is defined as $F(l) = \Pr[L \leq l]$. The corresponding marginal distribution function for capital income is denoted by $G(k) = \Pr[K \leq k]$. By Sklar's Theorem, there exists a copula function C_X such that for all x in $\mathbb{R}^2$ (Nelsen, 2006)

$$H_X(l, k) = C_X\{F(l), G(k)\} \quad (2.6)$$

The density of the joint distribution is obtained by differentiating with respect to l and k

$$h(l, k) = f(l)g(k)C_{FG}\{F(l), G(k)\} \quad (2.7)$$

where $C_{FG}\{F(l), G(k)\}$ is the copula density. The joint density can thus be expressed as the product of the marginals and the copula function. In the case of independence, the joint density is simply the product of the marginals. Because I am interested in top incomes, it makes sense to consider the *survival* copula. Survival analysis studies the probability of someone surviving beyond a certain point in time. Similarly, my analysis is concerned with the probability of having income above a certain level. As shown by Nelsen (2006), I can write the joint survival function very similar to the results above

$$\widehat{H_X}(l, k) = \Pr[L > l \text{ and } K > k] = \widehat{C_X}\{\widehat{F}(l), \widehat{G}(k)\} \quad (2.8)$$

where $\widehat{C_X}$ is the survival copula. The survival distributions (or complementary cumulative distribution functions) are given by $\widehat{F}(l) = 1 - F(l) = \Pr[L > l]$ and $\widehat{G}(k) = 1 - G(k) = \Pr[K > k]$.

The copula could be estimated by adopting a particular functional form, of which

there are many to choose from (Nelsen, 2006). In addition, one would have to estimate the marginal distributions which could be done either parametrically or non-parametrically (Trivedi and Zimmer, 2005). As I will show below, the joint distribution is asymmetric in the following sense: Those individuals at the top of the labour distribution are more likely to be at the top of the capital distribution compared with persons at the top of the capital distribution being at the top of the labour distribution.

It is difficult to find a functional form of the copula which could accommodate this asymmetry, which would also typically be ignored by a correlation coefficient. In their copula-based analysis of the joint distribution of income and wealth, Jäntti et al. (2012) argue that a more complex, asymmetric dependency could be captured by considering a mixture copula, i.e. combining different parametric copulas. However, they fail to fit the mixture model using their maximum likelihood approach. Overall, none of their four functional forms of the copula function can reproduce the correlation measures observed in the data. On the other hand, the marginal distributions are approximated reasonably well. They leave the search for a functional form, which captures the complexity of the dependency structure, for future research.

2.5.1 Association matrices as a discrete approximation

Given the difficulties in capturing the observed patterns of association, I have decided to proceed non-parametrically by considering association matrices. In studies of economic mobility, these are typically referred to as transition matrices (e.g. Atkinson, 1981). An association matrix is a discrete approximation to the copula density (Bonhomme and Robin, 2009). The copula for a discrete distribution is simply the cumulated association matrix (Dardanoni and Lambert, 2001). Furthermore, cumulating from above rather than below, yields the survival copula for a discrete distribution, which I refer to as the “inverse cumulative association matrix”.

Table 2.4 shows the association matrix for 1980. The factor incomes are broken down into eight unequally-spaced quantile groups: The bottom 50% ($\leq$ P50), the next 10% (P50-P60), the next 20% (P60-P80), the next 10% (P80-P90), the next 5% (P90-

P95), the next 4% (P95-P99), the next 0.5% (P99-P99.5) and the top 0.5% (>P99.5). The fractile categories always include the upper and exclude the lower cut-off, e.g. the P50-P60 group includes all observations which are greater than the median, and less than or equal to the 60th percentile. In the construction of percentiles, ties are broken by taking averages.⁶⁰

The literature on transition matrices discusses floor and ceiling effects (Hart, 1976). The diagonal entries at the bottom and the top (i.e. the top left and bottom right corners) overstate immobility (relative to the rest of the transition matrix), because individuals can only move in one direction (Atkinson et al., 1983). A similar issue applies to the association matrices in this chapter: Those observations in the top 0.5% (bottom 50%) of the labour distribution can only move downwards (upwards) or stay in the same position in terms of capital. However, I would argue that choosing very small fractiles at the top limits the extent of this problem (at least at the top of the distribution). Furthermore, my analysis does not focus on the diagonal elements, in contrast to studies of transition matrices, in which the diagonal has a particular interpretation (e.g. inter-generational persistence).

The association matrix in 2000 (Table 2.6) shows that 0.13% of observations were both in the top 0.5% of the labour and the top 0.5% of the capital income distribution, higher than 0.0025% (the frequency under the independence case, see Table 2.3), but less than the frequency under perfect association which would be 0.5% (Table 2.2).⁶¹ In other words, 26% (60%) of the top 0.5% of the labour distribution were in the top 0.5% (5%) of the capital distribution; more than 70% are in the top 10% of capital incomes.⁶² The association appears weaker for the top capital recipients: 52% (58%) of the top 0.5% of capital were in the top 5% (10%) of labour. Hence the association

⁶⁰I use the default method of Stata's `_pctile` command which takes averages where the empirical distribution function is flat.

⁶¹The matrix under perfect association is straightforward to obtain, it has all the weight on the diagonal. If labour and capital incomes are independent, the joint density equals the product of the marginals, which can be used to obtain the figures in Table 2.3. For example under independence, $\Pr(\text{bottom } 50\% \text{ labour} \cap \text{top } 0.5\% \text{ capital}) = \Pr(\text{bottom } 50\% \text{ labour}) \times \Pr(\text{top } 0.5\% \text{ capital}) = 0.5 \times 0.005 = 0.0025\%$

⁶²This is derived from the definition of a conditional probability and adding up the frequencies in the association matrix: $\Pr(\text{top } 5\% \text{ capital} \mid \text{top } 0.5\% \text{ labour}) = \Pr(\text{top } 5\% \text{ capital} \cap \text{top } 0.5\% \text{ labour}) / \Pr(\text{top } 0.5\% \text{ labour}) = (0.13\% + 0.05\% + 0.12\%) / 0.5\% = 60\%$

of capital and labour incomes is asymmetric.

Table 2.7 shows the “inverse cumulative association matrix” for 1980, which is obtained by adding up the entries in Table 2.4 starting in the bottom right-hand corner. Figure 2.8 offers another illustration of the asymmetry in the association between labour and capital. The line labelled “capital 1980” shows the cumulative frequency of different top quantiles of the labour distribution for the top 1% of capital recipients, i.e. going upwards along the column “top 1% of capital” in Table 2.7. Similarly, the line labelled “labour 2000” shows the cumulative frequencies of capital for the top 1% of labour recipients, i.e. going across the row “top 1% of labour” in Table 2.9. Figures 2.9 and 2.10 repeat this for the top 5% and top 0.5% respectively. In all three figures, the line of no association is obtained by cumulating the association matrix for the independence case (Table 2.3). Between 1980 and 2000, all three figures show a continuous movement away from the line of no association, i.e. the two sources of income have become more strongly associated. The shift between 1980 and 1990 appears stronger for the top 1% (and top 0.5%) than for the top 5%. Again, the asymmetric association is apparent from the fact that for any particular year, the “labour” line lies above the “capital” line.

Table 2.10 displays the same information as Figure 2.8. It is the conditional distribution, as defined in footnote 62.⁶³ It shows the strong degree of association between labour and capital at the top: For example, a tax unit in the top 1% of labour incomes in 2000 is almost guaranteed (with 80% probability) to be in the top 20% of capital incomes. The association is weaker for the top capital recipients, although still more than half of their top 1% are in the top 10% of labour incomes. At least for the top 1% shown in Table 2.10, the association appears to have become stronger over time.

⁶³In order to obtain the first column in Table 2.10, I firstly take the cumulative frequencies from Table 2.9 starting at cell (top 1%, top 1%) and going upwards along the column, and secondly divide by the marginal frequency, which in this case is simply 1%.

2.5.2 Differences between cumulative association matrices

In order to test formally whether the degree of association between labour and capital has increased, I now consider differences in the (inverse) cumulative association matrices. My approach is taken from Atkinson (1981) who applies the test to transition matrices in order to measure economic mobility. This is a first-order dominance test, and one might not be able to rank all years according to this criterion. Consider the following two association matrices A and A^*

$$\begin{array}{ccccc}
 & \cdots & j-1 & j & \cdots \\
 \vdots & \ddots & \vdots & \vdots & \ddots \\
 A = \begin{array}{ccccc} i-1 & \cdots & p_{i-1,j-1} & p_{i-1,j} & \cdots \\ i & \cdots & p_{i,j-1} & p_{i,j} & \cdots \\ \vdots & \ddots & \vdots & \vdots & \ddots \end{array} & \begin{array}{ccccc} \cdots & j-1 & j & \cdots \\ \vdots & \ddots & \vdots & \vdots & \ddots \\ i-1 & \cdots & p_{i-1,j-1} + \gamma & p_{i-1,j} - \gamma & \cdots \\ i & \cdots & p_{i,j-1} - \gamma & p_{i,j} + \gamma & \cdots \\ \vdots & \ddots & \vdots & \vdots & \ddots \end{array}
 \end{array}$$

where i and j are particular fractile groups (of labour and capital), $p_{i,j}$ is the frequency in the association matrix, and $\gamma > 0$. A^* is obtained from A by a diagonalising switch, which adds γ to the diagonal elements and subtracts it from the off-diagonal elements. This switch increases the weight on the diagonal and hence A^* exhibits stronger association between labour and capital. This transformation leaves the marginal distributions unchanged, because the marginal distribution of say labour is obtained by adding up over the columns along a row, so γ simply cancels out.

Let α be the inverse cumulative association matrix of A , which is defined as

$$\begin{array}{ccccc}
 & \cdots & \geq j-1 & \geq j & \geq j+1 & \cdots \\
 \vdots & \ddots & \vdots & \vdots & \vdots & \ddots \\
 \alpha = \begin{array}{ccccc} \geq i-1 & \cdots & \alpha_{i-1,j} + \alpha_{i,j-1} - \alpha_{i,j} + p_{i-1,j-1} & \alpha_{i-1,j+1} + \alpha_{i,j} - \alpha_{i,j+1} + p_{i-1,j} & \alpha_{i-1,j+1} & \cdots \\ \geq i & \cdots & \alpha_{i,j} + \alpha_{i+1,j-1} - \alpha_{i+1,j} + p_{i,j-1} & \alpha_{i,j+1} + \alpha_{i+1,j} - \alpha_{i+1,j+1} + p_{i,j} & \alpha_{i,j+1} & \cdots \\ \geq i+1 & \cdots & \alpha_{i+1,j-1} & \alpha_{i+1,j} & \alpha_{i+1,j+1} & \cdots \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots \end{array}
 \end{array}$$

where $\alpha_{i,j} = \Pr(l > l_i \cap k > k_j)$. I have used the fact that $\alpha_{i,j} = \alpha_{i,j+1} + \alpha_{i+1,j} - \alpha_{i+1,j+1} + p_{i,j}$; the diagonal element $\alpha_{i+1,j+1}$ needs to be subtracted because adding the adjacent elements $\alpha_{i,j+1}$ and $\alpha_{i+1,j}$ double-counts these cells. All other cells follow from the same formula.

The inverse cumulative association matrix of A^* is denoted by α^* and is defined similarly as

$$\alpha^* = \begin{array}{cccccc} & \cdots & \geq j-1 & & \geq j & \geq j+1 & \cdots \\ & \vdots & \vdots & & \vdots & \vdots & \ddots \\ \geq i-1 & \cdots & \alpha_{i-1,j} + \alpha_{i,j-1} - (\alpha_{i,j} + \gamma) + p_{i-1,j-1} + \gamma & \alpha_{i-1,j+1} + (\alpha_{i,j} + \gamma) - \alpha_{i,j+1} + p_{i-1,j} - \gamma & \alpha_{i-1,j+1} & \cdots \\ \geq i & \cdots & (\alpha_{i,j} + \gamma) + \alpha_{i+1,j-1} - \alpha_{i+1,j} + p_{i,j-1} - \gamma & \alpha_{i,j+1} + \alpha_{i+1,j} - \alpha_{i+1,j+1} + p_{i,j} + \gamma & \alpha_{i,j+1} & \cdots \\ \geq i+1 & \cdots & \alpha_{i+1,j-1} & \alpha_{i+1,j} & \alpha_{i+1,j+1} & \cdots \\ & \vdots & \vdots & \vdots & \vdots & \ddots \end{array}$$

where I already used the following results: $\alpha_{i,j}^* = \alpha_{i,j} + \gamma$; $\alpha_{i,j-1}^* = \alpha_{i,j-1}$ and $\alpha_{i-1,j}^* = \alpha_{i-1,j}$ because γ cancels out. After canceling out γ , it is clear that the only difference between α and α^* is $\alpha_{i,j}^*$, such that $\alpha_{i,j}^* = \alpha_{i,j} + \gamma$.

Therefore, taking the difference between α^* and α yields the following result

$$\alpha^* - \alpha = \begin{array}{cccccc} & \cdots & \geq j-1 & \geq j & \geq j+1 & \cdots \\ & \vdots & \vdots & \vdots & \vdots & \ddots \\ \geq i-1 & \cdots & 0 & 0 & 0 & \cdots \\ \geq i & \cdots & 0 & \gamma & 0 & \cdots \\ \geq i+1 & \cdots & 0 & 0 & 0 & \cdots \\ & \vdots & \vdots & \vdots & \vdots & \ddots \end{array}$$

This means that if the difference between the inverse cumulative association matrices in years $t+1$ and t is everywhere positive, the association between labour and capital has increased between those years. In other words, the distribution in $t+1$ is closer to the positive diagonal and further away from a class model, where one class is at the top of the labour distribution and the other at the top of the capital distribution. My exposition, which follows Atkinson (1981), relies on association matrices. Decancq (2009) derives dominance criteria directly for (survival) copula functions. If

the survival copula in year $t + 1$ is everywhere greater than the survival copula in year t , dependence has increased.⁶⁴

Tables 2.11, 2.12, and 2.13 show the results from comparing the inverse cumulative association matrices for the years 2000, 1990 and 1980. In Table 2.11, a positive number means that the inverse cumulative association matrix (or survival copula) in 2000 exceeds the 1980 level. If all entries were positive, dependence between labour and capital would have increased unambiguously. This does not happen in any of the three comparisons.⁶⁵ However, when comparing 2000 to 1980 (Table 2.11) there are few negative entries, so it appears as if dependence increased, at least for the top 20%. One might draw the same conclusion from the comparison of 2000 with 1990 (Table 2.13), ignoring the very small difference at the very top of the distribution. Moreover, ignoring the small difference of less than 0.1, between 1990 and 2000 the association increased in the whole distribution, not just at the top. On the other hand, the ranking of 1990 and 1980 appears inconclusive (Table 2.12). This means that most of the change between 2000 and 1980 occurred after 1990, which might suggest that the TRA86 had a limited effect on the association.⁶⁶ By allocating all of S-Corporation income to capital, I might have successfully limited some of the most pervasive income switching.

⁶⁴That is, dependence according to the upper orthant dependence relation, as defined in Decancq (2009) has increased.

⁶⁵Aaberge et al. (2013) propose a test of second-order dominance, which involves integrating the copula function. This is still very preliminary. Nelsen (2006) shows that the rescaled volume under the copula function is the Spearman's rho rank correlation measure. Atkinson (1981) suggests a discrete variant of this approach applied to measuring mobility. However, if one ignores very small differences, the results presented here suggest first-order dominance, i.e. the association of capital and labour incomes increased over this 20-year period.

⁶⁶We can see the same results from the cross-distribution charts (Figures 2.8, 2.9, 2.10). Tables 2.11, 2.12, 2.13 are simply the differences between the respective lines. For example, the share of tax units which are both in the top 5% of labour and top 40% of capital declined by 0.02pp between 1990 and 1980 (as shown by the "labour 1990" and "labour 1980" lines in Figure 2.9).

2.6 Conclusion

This chapter has contributed to the study of top income shares in two ways. First, I showed that the evolution of top income shares is robust to controlling for tax unit size. Second, I found that the association between labour and capital incomes has become stronger since the 1980s (at least at the very top) and has contributed to the increasing top income share. I used a copula framework to estimate a measure of association that is independent of changes in the marginal distributions, in contrast to the association measures in factor income decompositions. This association is found to be asymmetric, as top wage earners are more likely to be at the top of the capital distribution compared with top capital income recipients being at the top of the wage distribution.

Although it concerns by definition only a small number of people, the study of top incomes is important for at least four reasons. First of all, they account for a substantial share of national income (Atkinson et al., 2011), e.g. the top 1% receive close to 20% of gross income in 2010. In a progressive tax system they pay an even larger share of total taxes, so the evolution of their incomes has significant fiscal implications. Second, as a result of these economic resources, it is probably reasonable to expect that the super-rich have disproportionate control over the democratic political process (Solow, 2014). Third, it matters how these incomes were generated and whether they reflect inefficient rent capture by particular sectors or occupations. Fourth, top income shares are important for measuring inequality accurately, because many standard measures of inequality are less sensitive to changes in the tails where substantial shifts occurred in recent years.

This chapter has contributed to the study of the functional distribution of income. For the classical economists, the mapping between the functional and the personal income distributions was clear: Landlords were rich, workers were poor, and capitalists in between. This relationship has become more complex for two reasons (Atkinson, 2009): First, many people have more than one source of income. My analysis has shown that at least at the top of the distribution, the association between capital and labour has become even stronger. Second, there exists substantial inequality within

factor incomes. Section 2.4 showed that the inequality within the capital and labour distributions has increased over time.

This chapter has two immediate policy implications. First, I showed that the inequality in labour incomes increased substantially, and in relative terms even more than the inequality in capital incomes. This may be explained by market-based factors, such as a superstar theory together with a globalised economy (Atkinson, 2008) or the spread of winner-take-all pay structures (Gordon and Dew-Becker, 2008). Some individuals at the top of the income distribution have considerable freedom to set their own compensation, so the process by which their pay is determined is perhaps best modelled in a job-matching model with Nash bargaining (Alvaredo et al., 2013a). For example, this could be the CEO and his compensation committee. Governments influence this bargaining process through the top marginal tax rate which affects the reward to bargaining more aggressively. Governments can also exert a certain amount of moral suasion (Atkinson, 2013b) and try to shape social norms (Piketty and Saez, 2007b). Furthermore, to understand top earnings it is crucial to explain the vast and increasing rents in some professions such as the financial sector (Kaplan and Rauh, 2010). Legislation, for example financial sector regulation, plays a very important role.

The second policy implication concerns the association between capital and labour incomes. The analysis has shown that top earners are increasingly at the top of the distribution of capital incomes. These incomes are the returns to a stock of capital that has been accumulated during a period when top income tax rates were low compared with both US history and other countries in the world (Alvaredo et al., 2013a). For example, the top marginal tax rate in the US declined from 91% in the early 1960s to 28% in 1986. If fiscal policy continues in a similar vein, the association between capital and labour income is likely to increase further.

2.7 Main Tables and Figures

Figure 2.1: Income share of the Top 1% of income recipients




Figure 2.2: Income composition of the Top 1% of income recipients




Table 2.1: Summary statistics

	1970	1980	1986	1987	1990	2000	2005
No. observations	95,316	171,508	75,400	109,939	89,452	145,663	132,766
No. of tax units	74,286,993	93,902,125	103,037,255	106,996,252	113,787,293	129,373,494	134,433,082
Fraction of filers	0.93	0.94	0.93	0.95	0.96	0.96	0.92
A. Average and threshold incomes (<i>excluding</i> capital gains) per tax unit (in 2005 USD)							
Overall	37,037	36,938	37,851	38,902	40,362	48,227	46,239
P90-95: average	82,496	86,398	92,319	93,289	95,275	112,726	111,270
P90-95: threshold	74,509	77,747	82,348	83,067	83,997	97,233	95,640
P95-99: average	117,023	121,115	130,708	134,713	143,294	183,011	180,749
P95-99: threshold	93,291	97,767	105,506	107,000	110,206	133,716	132,316
Top 1%: average	288,927	304,479	349,375	423,314	529,275	800,489	818,507
Top 1%: threshold	171,694	175,949	185,998	202,419	224,892	302,118	302,937
B. Average and threshold incomes (<i>including</i> capital gains) per tax unit (in 2005 USD)							
Overall	38,107	38,663	42,822	40,944	41,771	53,515	50,794
P90-95: average	83,627	88,546	97,311	95,718	96,768	118,235	115,666
P90-95: threshold	75,307	79,426	85,825	84,828	85,025	101,042	98,553
P95-99: average	120,817	128,080	148,189	142,856	148,202	200,263	196,386
P95-99: threshold	95,118	100,758	113,027	110,176	112,450	141,479	138,600
Top 1%: average	342,520	387,390	651,571	522,301	603,468	1,132,188	1,106,251
Top 1%: threshold	182,400	194,297	236,814	225,546	238,873	346,489	346,670
C. Tax unit size adjustment							
Avg. tax unit size	2.48	2.23	N/A	N/A	N/A	2.02	2.04
Fraction dependent tax units	0	0	N/A	N/A	N/A	0.09	0.07
Dependents: Avg. excl. KG			N/A	N/A	N/A	4,726	5,408
Dependents: Avg. incl. KG			N/A	N/A	N/A	5,613	5,766
D. Summary statistics for observations dropped in Part II							
Frac. neg. labour	0	0	0	0	0	0	0
Frac. neg. business	0.04	0.05	0.06	0.06	0.05	0.05	0.05
Frac. neg. capital (incl. KG)	0.04	0.04	0.04	0.04	0.04	0.05	0.07
Frac. dropped	0.07	0.08	0.10	0.10	0.09	0.09	0.11
Dropped obs.: Avg. excl. KG	53,618	57,480	61,070	67,449	67,959	73,441	66,781
Dropped obs.: Avg. incl. KG	56,640	65,006	83,102	76,195	74,016	89,210	80,899
No. observations Part II	71,038	121,039	46,439	81,185	67,133	112,707	99,198
No. of tax units Part II	68,864,271	86,323,526	92,516,221	96,828,263	103,481,852	117,531,939	119,112,529
E. Share of other income in total income (including capital gains): Income excluded in Part II							
Overall	0.00	0.00	0.00	0.00	0.00	0.00	-0.01
P90-P95	0.00	0.01	0.01	0.01	0.01	0.01	0.01
P95-P99	0.00	0.01	0.01	0.01	0.01	0.01	0.01
Top 1%	0.01	0.00	0.00	0.01	0.02	0.01	0.01

See the description in Section 2.2.5. All average and threshold incomes are expressed in 2005 USD (CPI from Saez, 2013). The overall average income (incl. and excl. KG) includes the income imputed for the non-filers.

Figure 2.3: Accounting for tax unit size (Top 0.5% and Top 1%)




Figure 2.4: Accounting for tax unit size (Top 5% and Top 10%)




Figure 2.5: Factor income decomposition of Top 1% share




Figure 2.6: Factor income decomposition of Top 5% share




Figure 2.7: Factor income decomposition of Top 0.5% share




Table 2.2: Association matrix: Perfect association (frequencies in %)

Perf. Assoc.	Capital								
Labour	$\leq P50$	$P50-P60$	$P60-P80$	$P80-P90$	$P90-P95$	$P95-P99$	$P99-P99.5$	$> P99.5$	Total
$\leq P50$	50	0	0	0	0	0	0	0	50.00
$P50-P60$	0	10	0	0	0	0	0	0	10.00
$P60-P80$	0	0	20	0	0	0	0	0	20.00
$P80-P90$	0	0	0	10	0	0	0	0	10.00
$P90-P95$	0	0	0	0	5	0	0	0	5.00
$P95-P99$	0	0	0	0	0	4	0	0	4.00
$P99-P99.5$	0	0	0	0	0	0	0.5	0	0.50
$> P99.5$	0	0	0	0	0	0	0	0.5	0.50
Total	50.00	10.00	20.00	10.00	5.00	4.00	0.50	0.50	100.00

Table 2.3: Association matrix: Independence (frequencies in %)

Independ.	Capital								
Labour	$\leq P50$	$P50-P60$	$P60-P80$	$P80-P90$	$P90-P95$	$P95-P99$	$P99-P99.5$	$> P99.5$	Total
$\leq P50$	25	5	10	5	2.5	2	0.25	0.25	50.00
$P50-P60$	5	1	2	1	0.5	0.4	0.05	0.05	10.00
$P60-P80$	10	2	4	2	1	0.8	0.1	0.1	20.00
$P80-P90$	5	1	2	1	0.5	0.4	0.05	0.05	10.00
$P90-P95$	2.5	0.5	1	0.5	0.25	0.2	0.025	0.025	5.00
$P95-P99$	2	0.4	0.8	0.4	0.2	0.16	0.02	0.02	4.00
$P99-P99.5$	0.25	0.05	0.1	0.05	0.025	0.02	0.0025	0.0025	0.50
$> P99.5$	0.25	0.05	0.1	0.05	0.025	0.02	0.0025	0.0025	0.50
Total	50.00	10.00	20.00	10.00	5.00	4.00	0.50	0.50	100.00

Table 2.4: Association matrix: 1980 (frequencies in %)

1980	Capital								
Labour	$\leq P50$	$P50-P60$	$P60-P80$	$P80-P90$	$P90-P95$	$P95-P99$	$P99-P99.5$	$> P99.5$	Total
$\leq P50$	30.14	3.48	6.32	4.87	2.62	2.09	0.26	0.23	50.00
$P50-P60$	5.63	1.05	1.83	0.74	0.46	0.24	0.02	0.02	10.00
$P60-P80$	9.26	2.67	4.88	1.72	0.83	0.54	0.05	0.04	20.00
$P80-P90$	3.25	1.59	3.18	1.11	0.42	0.40	0.03	0.03	10.00
$P90-P95$	1.15	0.72	1.94	0.65	0.25	0.24	0.02	0.02	5.00
$P95-P99$	0.54	0.45	1.60	0.70	0.30	0.32	0.05	0.05	4.00
$P99-P99.5$	0.03	0.02	0.15	0.11	0.06	0.08	0.02	0.02	0.50
$> P99.5$	0.01	0.01	0.10	0.09	0.06	0.10	0.04	0.08	0.50
Total	50.01	10.00	19.99	10.00	5.00	4.00	0.50	0.50	100.00

Reading: 30.14% of taxpayers are both in the bottom half of the labour distribution and the bottom half of the capital distribution. This means, 60.28% of those taxpayers who are in the bottom half of labour incomes, are also in the bottom half of the capital distribution.

Table 2.5: Association matrix: 1990 (frequencies in %)

1990	Capital								
Labour	$\leq P50$	$P50-P60$	$P60-P80$	$P80-P90$	$P90-P95$	$P95-P99$	$P99-P99.5$	$> P99.5$	Total
$\leq P50$	29.21	2.97	7.41	5.23	2.68	2.11	0.22	0.17	50.00
$P50-P60$	6.11	0.83	1.61	0.68	0.50	0.21	0.02	0.02	10.00
$P60-P80$	9.79	2.70	4.36	1.64	0.86	0.56	0.06	0.04	20.00
$P80-P90$	3.31	1.87	2.99	1.02	0.38	0.37	0.03	0.03	10.01
$P90-P95$	1.12	0.98	1.82	0.57	0.21	0.24	0.02	0.02	4.99
$P95-P99$	0.44	0.59	1.59	0.66	0.26	0.35	0.06	0.05	4.00
$P99-P99.5$	0.02	0.03	0.14	0.11	0.05	0.08	0.05	0.03	0.50
$> P99.5$	0.01	0.01	0.07	0.08	0.06	0.09	0.04	0.14	0.50
Total	50.02	9.98	20.00	10.00	5.00	4.00	0.50	0.50	100.00

Table 2.6: Association matrix: 2000 (frequencies in %)

2000	Capital								
Labour	$\leq P50$	$P50-P60$	$P60-P80$	$P80-P90$	$P90-P95$	$P95-P99$	$P99-P99.5$	$> P99.5$	Total
$\leq P50$	30.09	2.90	8.13	4.90	2.18	1.54	0.14	0.11	50.00
$P50-P60$	6.10	0.96	1.53	0.70	0.43	0.24	0.02	0.01	10.00
$P60-P80$	9.61	2.77	4.29	1.73	0.95	0.56	0.06	0.04	20.00
$P80-P90$	3.13	1.83	2.88	1.02	0.58	0.47	0.05	0.03	10.00
$P90-P95$	0.85	0.97	1.65	0.74	0.34	0.39	0.04	0.03	5.01
$P95-P99$	0.31	0.45	1.35	0.74	0.39	0.57	0.09	0.09	3.99
$P99-P99.5$	0.01	0.02	0.10	0.10	0.06	0.11	0.05	0.04	0.50
$> P99.5$	0.00	0.01	0.05	0.07	0.07	0.12	0.05	0.13	0.50
Total	50.10	9.91	19.99	10.00	5.00	4.00	0.50	0.50	100.00

Table 2.7: Inverse cumulative association matrix: 1980 (frequencies in %)

1980	Capital							
Labour	Whole	Top 50%	Top 40%	Top 20%	Top 10%	Top 5%	Top 1%	Top 0.5%
Whole	100.00	49.99	39.99	20.00	10.00	5.00	1.00	0.50
Top 50%	50.00	30.12	23.60	9.93	4.81	2.42	0.51	0.27
Top 40%	40.00	25.75	20.29	8.45	4.07	2.15	0.47	0.25
Top 20%	20.00	15.01	12.22	5.26	2.60	1.52	0.38	0.20
Top 10%	10.00	8.26	7.06	3.28	1.73	1.06	0.32	0.17
Top 5%	5.00	4.41	3.93	2.09	1.19	0.77	0.27	0.15
Top 1%	1.00	0.95	0.92	0.68	0.47	0.36	0.17	0.10
Top 0.5%	0.50	0.49	0.48	0.38	0.29	0.22	0.12	0.08

This table is obtained by adding up the frequencies in Table 2.4 from the top of the distribution.

Reading: 23.60% of observations are both in the top 50% of labour and the top 40% of capital incomes.

Table 2.8: Inverse cumulative association matrix: 1990 (frequencies in %)

1990	Capital							
Labour	Whole	Top 50%	Top 40%	Top 20%	Top 10%	Top 5%	Top 1%	Top 0.5%
Whole	100.00	49.98	40.00	20.00	9.99	5.00	1.00	0.50
Top 50%	50.00	29.19	22.18	9.59	4.82	2.50	0.61	0.33
Top 40%	40.00	25.31	19.13	8.15	4.06	2.25	0.57	0.31
Top 20%	20.00	15.10	11.61	5.00	2.55	1.60	0.47	0.27
Top 10%	9.99	8.41	6.80	3.17	1.75	1.18	0.42	0.24
Top 5%	5.00	4.54	3.91	2.10	1.25	0.89	0.37	0.22
Top 1%	1.00	0.97	0.94	0.72	0.53	0.43	0.26	0.17
Top 0.5%	0.50	0.49	0.49	0.41	0.33	0.27	0.18	0.14

Table 2.9: Inverse cumulative association matrix: 2000 (frequencies in %)

2000	Capital							
Labour	Whole	Top 50%	Top 40%	Top 20%	Top 10%	Top 5%	Top 1%	Top 0.5%
Whole	100.00	49.90	39.99	20.00	10.00	5.00	1.00	0.50
Top 50%	50.00	29.98	22.98	11.11	6.02	3.20	0.75	0.39
Top 40%	40.00	26.09	20.04	9.71	5.32	2.93	0.71	0.37
Top 20%	20.00	15.69	12.42	6.37	3.71	2.27	0.61	0.33
Top 10%	10.00	8.82	7.38	4.22	2.57	1.71	0.52	0.29
Top 5%	4.99	4.67	4.19	2.68	1.78	1.26	0.45	0.26
Top 1%	1.00	0.98	0.96	0.80	0.63	0.50	0.27	0.17
Top 0.5%	0.50	0.50	0.49	0.44	0.37	0.30	0.18	0.13

Figure 2.8: Depiction of the inverse cumulative association matrix: Top 1%




Figure 2.9: Depiction of the inverse cumulative association matrix: Top 5%




Figure 2.10: Depiction of the inverse cumulative association matrix: Top 0.5%




Table 2.10: Conditional distribution for Top 1% (frequencies in %)

	Capital											
	Top 1%			Top 5%			Top 10%			Top 20%		
Labour	2000	1990	1980	2000	1990	1980	2000	1990	1980	2000	1990	1980
Top 1%	27	26	17	50	43	36	63	53	47	80	72	68
Top 5%	45	37	27									
Top 10%	52	42	32									
Top 20%	61	47	38									

Reading: Of those tax units in the top 1% of capital income, 52% were in the top 10% of labour income in 2000 (compared with 42% in 1990 and 32% in 1980).

Table 2.11: Change in inverse cumulative association matrix: 2000 and 1980 (in pp)

2000 - 1980	Capital							
Labour	Whole	Top 50%	Top 40%	Top 20%	Top 10%	Top 5%	Top 1%	Top 0.5%
Whole	0.00	-0.08	0.00	0.00	0.00	0.00	0.00	0.00
Top 50%	0.00	-0.13	-0.63	1.19	1.21	0.78	0.24	0.12
Top 40%	0.00	0.34	-0.25	1.27	1.25	0.78	0.24	0.12
Top 20%	0.00	0.68	0.20	1.11	1.10	0.75	0.23	0.12
Top 10%	0.00	0.56	0.32	0.94	0.85	0.65	0.21	0.12
Top 5%	-0.01	0.26	0.26	0.59	0.59	0.49	0.18	0.11
Top 1%	0.00	0.03	0.04	0.13	0.16	0.15	0.09	0.06
Top 0.5%	0.00	0.01	0.01	0.06	0.08	0.08	0.05	0.05

Reading: The share of tax units who were both in the top 0.5% of labour and top 0.5% of capital incomes was 0.05pp greater in 2000 compared with 1980. Negative cells are highlighted in red.

Table 2.12: Change in inverse cumulative association matrix: 1990 and 1980 (in pp)

1990 - 1980	Capital							
Labour	Whole	Top 50%	Top 40%	Top 20%	Top 10%	Top 5%	Top 1%	Top 0.5%
Whole	0.00	0.00	0.01	0.00	-0.01	0.00	0.00	0.00
Top 50%	0.00	-0.93	-1.43	-0.34	0.01	0.08	0.10	0.06
Top 40%	0.00	-0.44	-1.16	-0.30	-0.01	0.11	0.10	0.07
Top 20%	0.00	0.08	-0.61	-0.27	-0.05	0.09	0.10	0.07
Top 10%	-0.01	0.14	-0.27	-0.11	0.02	0.11	0.10	0.07
Top 5%	0.00	0.13	-0.02	0.01	0.06	0.11	0.10	0.07
Top 1%	0.00	0.02	0.02	0.05	0.06	0.07	0.08	0.07
Top 0.5%	0.00	0.01	0.01	0.03	0.04	0.05	0.06	0.06

Table 2.13: Change in inverse cumulative association matrix: 2000 and 1990 (in pp)

2000 - 1990	Capital							
Labour	Whole	Top 50%	Top 40%	Top 20%	Top 10%	Top 5%	Top 1%	Top 0.5%
Whole	0.00	-0.08	-0.01	0.00	0.01	0.00	0.00	0.00
Top 50%	0.00	0.79	0.80	1.53	1.20	0.70	0.13	0.05
Top 40%	0.00	0.78	0.92	1.56	1.26	0.68	0.14	0.06
Top 20%	0.00	0.60	0.80	1.38	1.15	0.67	0.13	0.05
Top 10%	0.01	0.42	0.58	1.05	0.82	0.54	0.11	0.05
Top 5%	-0.01	0.13	0.28	0.58	0.53	0.37	0.08	0.04
Top 1%	0.00	0.01	0.02	0.08	0.10	0.08	0.01	0.00
Top 0.5%	0.00	0.00	0.00	0.02	0.04	0.03	0.00	-0.01

2.A Appendix

In this section, I review the literature on a number of potential issues which could bias the results obtained from tax record data. Firstly, I address the issue of tax evasion and the effects of changes in the tax code. In the US, the top marginal tax rate declined from 91% in early 1960s to 28% in 1986. By definition, it is difficult to measure the extent of tax avoidance. However, top shares excluding and including capital gains, an important channel of tax avoidance, moved very similarly (Piketty et al., 2011). In addition, the increase in top income shares is confirmed by independent evidence from salaries. Tax reporting behaviour might explain short-run fluctuations but cannot account for the long-term trends (Alvaredo et al., 2013a). Secondly, I discuss how the income definition used in this chapter relates to a Haig-Simons definition and why I use before-tax rather than after-tax incomes.

2.A.1 Tax evasion

If the rich are more likely to evade taxes, then the results are a lower bound on the level of the top income share. For tax evasion to explain the increase in the top income share after WWII, one would require that tax evasion has increased, and that it increased faster at the top (Atkinson et al., 2011). Tax evasion is hard to measure because (1) respondents are unlikely to report truthfully; (2) there exists a fine line between tax evasion and (legal) avoidance; (3) some misstatement of taxes is unintentional (Slemrod, 2007).

Using information from random intensive audits, the IRS estimated net underreporting to be 13.7% of the tax liability in 2001 (Slemrod, 2007). Slemrod further argues that non-farm business income is most prone to tax evasion because income such as wages, interest and dividends are reported by those who pay them out. In addition, wages are subject to tax withholding by the employer. There is some evidence to support this argument. Conditional on other characteristics, the self-employed have a higher propensity to consume food in the UK (Pissarides and Weber, 1989) and higher charitable donations relative to their (reported) income in the USA (Feldman

and Slemrod, 2007). Assuming that food consumption and donations are not related to the main source of income, this would suggest that the reported self-employed incomes are understated. On the other hand, Alvaredo and Pisano (2010) find that the difference in top income shares between Italy and Anglo-Saxon countries cannot be explained by tax evasion of the self-employed and small businesses.

The more important question is whether the extent of tax evasion varies along the income distribution. Christian (1994) shows that relative to their true income the rich evade less than the poor. This evidence might not be fully conclusive because he uses data from the IRS's random audits which might miss some non-compliance and ignore business incomes (Slemrod, 2007). Using evidence from a tax amnesty in Argentina, evasion seems to be lowest at the bottom and the top of the income distribution (Alvaredo, 2010). This could be explained by hard-to-evade wage income being prominent at the bottom and enforcement being more stringent at the top, something that is also confirmed by Leigh and van der Eng (2009) for Indonesia.

In sum, I agree with Atkinson et al. (2011) that for the purpose of measuring top incomes, legal tax-exemptions of capital incomes might be at least as important as outright tax evasion. While it is true that the richest have most to gain from evading taxes, the IRS also has most to gain from targeting enforcement on tax units with the highest liabilities (Atkinson et al., 2011). Although IRS enforcement might have been unreliable in the early years of income taxation, I have no reason to expect any variation in enforcement since the 1960s.

2.A.2 Changes in the tax code over time

Over the 40-year period under analysis, there have been a number of changes in the US tax code. I concentrate on the legislative changes which affected top incomes. In the early 1960s top marginal income tax rates declined from 91% to 70%. The Nixon administration's Tax Reform Act of 1969 introduced a 50% maximum tax on earned income. The Reagan administration Tax Reform Act of 1986 (TRA86) then reduced the top income tax rate to 28%, the lowest since the end of WWII. By 1993, top marginal

rates had risen again to 39.6% and in 2003 they were cut to 35%. The Clinton tax increase in 1993 increased the top personal income tax rate above corporation tax, which had remained unchanged since TRA86 at 34%. Since 1997, capital gains have again received a preferential tax treatment. Other tax changes since 1960 which mostly affected the middle of the distribution include the increase in payroll taxes financing social security and Medicare. The US federal tax system as a whole has become less progressive, which is largely due to the increase in payroll taxes and the fall in corporate and estate taxes (Piketty and Saez, 2007a).

While some of the tax cuts in the 1960s were large, Saez (2004) argues that, compared with TRA86, their impact on behaviour was limited, because top income tax rates remained above the corporate and capital gains tax rates. The TRA86 has been called “the most sweeping change in the US federal income tax since it became a tax that applied to most citizens during World War II” (Auerbach and Slemrod, 1997, p. 589) and I discuss it in more detail.

TRA86 consisted of a number of measures: First, as already mentioned, the top marginal income tax rate was reduced from 50% to 28%. Other marginal rates were also cut, but over 40% of tax units had their rates unchanged or increased (Hausman and Poterba, 1987). Marginal tax rates became non-monotonic with the highest bracket facing a lower tax rate than the previous bracket (Auerbach and Slemrod, 1997). Second, the law lowered the corporation tax from 46% to 34%. This was the first time in the USA that the top personal income tax rate was lower than the corporate tax rate (Fullerton, 1996).

Because TRA86 aimed to be revenue-neutral (Saez et al., 2012), the remaining provisions broadened the tax base to counteract the revenue shortfalls from the rate reductions. The third component of the TRA86 was the inclusion of all realised capital gains in taxable income. Previously, only 40% of long-term capital gains were subject to income tax. For someone in the top tax bracket, the effective capital gains tax rate rose from 20% (40% of capital gains taxed at 50%) to 28% (100% of gains taxed at 28%) (Slemrod, 1996). Lastly, the law removed a number of deductions and tax shel-

tering opportunities, particularly at the corporate level. Hence despite the reduction in the corporation tax, overall corporate tax revenues were projected to remain constant (Auerbach and Slemrod, 1997).

2.A.3 Effects of changes in the tax code

Slemrod (1992) suggests a useful hierarchy of behavioural responses to changes in taxation, ranked by responsiveness. The most responsive is the timing of incomes, such as the realisation of capital gains, which shifts income over time. Next are accounting responses, such as the shifting of income from the corporate to the personal tax base. Finally, there may exist real responses such as increased economic activity, which are hardest to detect. In my empirical analysis I consistently define income as gross income instead of taxable income thus taking care of the numerous changes in deductions and exemptions, following Piketty and Saez (2007b). In addition, I have adjusted capital gains to account for the above-mentioned changes in the tax base, which is explained in more detail in section 2.2.2.

The first issue of shifting income over time affects especially capital gains and stock options. Capital gains are taxed only when they are realised, which can be timed in response to changes in the tax code. Stock options are taxed as wage income when they are exercised, and the timing could depend on changes in the personal tax rate.⁶⁷ Once exercised, this income is treated like other capital income. Of course, the timing of stock options and capital gains is also affected by (expected) asset price movements.⁶⁸

For top income recipients, TRA86 increased taxes on capital gains and lowered the personal income tax, starting from tax year 1987. From Figure 2.2, we can clearly see

⁶⁷Stock options have certainly increased in importance between 1960 and 2005. Saez (2004) argues that stock options contributed to the substantial increase in the top wage income shares.

⁶⁸The importance of asset price movements is illustrated by the bust of the dot-com boom. Saez et al. (2012) report a spike in the wage share of the top 0.01% in 2000. This could be explained by employees in the new economy, where stock options were frequently used (Murphy, 2002), exercising their stock options before markets started to collapse. Figure 2.2 shows that the labour share excluding capital gains peaked in 2000, but it did not fall substantially until 2003, which was already well into the dot-com crash. Figure 2.2 also shows that the capital share including capital gains peaked in 1999 at 36.8% and then dropped substantially to 31.0% in 2001 and 28.1% in 2002. This could be explained by tax units cashing in on capital gains. On the other hand, the crash in 2001 and 2002 would reduce any capital gains mechanically.

the spike in the capital income share including capital gains in 1986: It increased from 45.8% in 1985 to 55.2% in 1986, falling to 34.4% in 1987. Overall, capital gains realisations increased by 96.3% in 1986 compared with 1985 (Slemrod, 1996).⁶⁹ The shifting of capital gains realisations also affects the overall share of total income (Figure 2.1): The income share excluding capital gains increases continuously between 1985 and 1987, whereas the income share including capital gains (in ranking and in the share) increases by more than 3pp in 1986 before dropping back to the 1985 level.

On the other hand, the anticipated decline in the income tax provided incentives to postpone the exercise of stock options, thus reducing the labour share. The labour share (excl. capital gains) is indeed lower in 1986 compared with 1987, but it is even lower in 1985, so in the absence of a true counterfactual, one cannot draw any definite conclusions. It might also be that stock options were not yet an important source of income in the mid-1980s.⁷⁰ Or it could be that stock options are not an important source for the top 1%: Saez et al. (2012) find a clear spike in the wage share of the top 0.01% in 1988. In 1993, top income tax rates increased from 31% to 39.6%, which prompted many executives to exercise their stock options in 1992 (Goolsbee, 2000) and there is some evidence that financial service institutions paid bonuses earlier than usual (Feenberg and Poterba, 2000). Figure 2.1 shows a peak of the top income share excluding capital gains in 1992.⁷¹ According to Figure 2.2, the labour share (excl. capital gains) jumps from 72.7% in 1991 to 77.3% in 1992, but it remains virtually unchanged after that. The wage share of the top 0.01%, where stock options might be more important, shows a clear peak in 1992 (Saez et al., 2012).

The next response in the Slemrod hierarchy is how changes in the tax code can lead to a relabelling of income from one source to the other. This would obviously change the composition of incomes, which is one of the topics of this thesis chapter. I

⁶⁹In Sweden, capital gains realisations were shifted from 1990 to 1991, when capital gains taxes fell substantially (Roine and Waldenström, 2012).

⁷⁰Murphy (2002) shows that the share of stock options (valued when granted using the Black-Scholes formula) in CEO compensation amongst the S&P 500 increased from 27% in 1992 to 51% in 2000.

⁷¹The top 1% income share excluding capital gains increased from 12.2% in 1991 to 13.5% in 1992, before falling to 12.8% in 1993 (Piketty and Saez, 2007b).

distinguish between two types of income shifting: First, the top income tax rate fell below the corporate rate for the first time in 1986 thus providing incentives to move income from the corporate to the personal income tax base. Second, TRA86 equalised taxes on earned income and capital gains, thus stopping the preferential treatment of capital gains.

In the US, corporate profits are subject to the corporate profit tax.⁷² However, partnerships, sole proprietorships, or otherwise unincorporated businesses, are taxed at the personal level. Businesses with few shareholders can choose to transfer to S-corporation status which is subject to the personal income tax (Saez et al., 2012).⁷³ The reduction of the personal tax below the corporate tax provided incentives for eligible businesses to switch to such pass-through entities (Auerbach and Slemrod, 1997). The number of filings to obtain S-corporation status increased substantially following TRA86 (Slemrod, 1996). An increasing share of S-corporation income was received by top income tax units. Amongst the top 0.01%, there was also a strong increase in S-corporation and partnership incomes between 1986 and 1988, the latter mostly driven by changes in legislation regarding partnership losses (Saez et al., 2012). It is, however, interesting that in 1993 when top personal income taxes rose to 39.6% (while the corporate rate was 35%) only few S-corporations switched back to C-corporation status.

The second major effect of TRA86 on the income composition concerns capital gains. The old tax schedule incentivised business owners to retain earnings in the corporate sector (taxed at the lower corporate rate), thus accumulating capital gains which received a preferential tax treatment. The increase in the capital gains and corporate income tax both contributed towards paying higher wages. In addition, for companies that retained C-corporation status, TRA86 removed the preferential treatment of capital gains over dividends.⁷⁴ In the medium-term, I would thus expect

⁷²If profits are paid out as dividends, they are taxed again as personal income. If profits are retained, they might create capital gains which would be taxed upon realisation.

⁷³There are further restrictions to an S-corporation (Saez, 2004): The number of shareholders is limited to currently 100. It can only issue one class of stock. It cannot be a subsidiary of another corporation.

⁷⁴This characterisation of the choice between capital gains and dividends might be too simplistic.

a decline in capital gains (Atkinson et al., 2011).⁷⁵ However, there exists mixed evidence on how much tax units can control the choice between capital gains and dividends.⁷⁶

I have reviewed a number of ways in which tax changes, particularly the TRA86, can affect top income shares and their composition. There are substantial short-run responses by top income recipients, but these effects mostly concern timing and relabelling, the first two tiers of the Slemrod (1992) hierarchy. Since my analysis covers a period of 45 years, I expect these short-term effects to be less important (Saez et al., 2012). Furthermore, I address some of the effects of the TRA86 on income composition by allocating S-corporation income entirely to capital income, which is also done by CBO (2012).⁷⁷ The increase in the wage share at the top appears too large and persistent to be accounted for purely by changes in relative taxation (Saez, 2004). It is also corroborated by independent evidence on executive compensation (Piketty and Saez, 2007b). Furthermore, a comparison with Canada suggests that the increase in top salaries in the USA is real and not simply the result of fiscal manipulation (Saez, 2004).⁷⁸ In addition, one could argue that retained earnings are paid out to individuals eventually as dividends or capital gains, so the shifts between corporate and personal incomes should only have a limited effect on top income shares (including capital gains)

There exists a substantial public finance literature trying to explain why corporations pay dividends instead of capital gains, which in most cases are taxed at a lower rate, although this is not the case for top income earners after TRA86 (Rosen and Gayer, 2010). These explanations include managerial control issues making dividends more attractive, and models in which dividends serve as a signal of financial soundness.

⁷⁵This is not readily apparent from Figures 2.1 and 2.2. Perhaps this is because of the one-off peak in capital gains realisations before the tax increases of TRA86 took effect. Furthermore, the magnitude of capital gains realisations also depends on other factors such as expected asset price movements.

⁷⁶During the early 1960s, dividends accounted for more than 60% of the incomes of the top 0.01%, although they faced a tax rate of 80% on average whereas long-term capital gains were taxed at 25% (Saez et al., 2012). This suggests that these tax units had little control over the choice of dividends or capital gains. This might be because they are outvoted by investors, such as pension funds and universities, which face a zero income tax. The 2003 dividend tax cut from 35% to 15% (for tax units in the top bracket) led to a strong increase in dividend payments, but not for corporations held by non-taxable institutions (Chetty and Saez, 2005).

⁷⁷This has been suggested by Emmanuel Saez. I split total income into labour and capital income. Business income is allocated two-thirds to labour and one-third to capital. Because of the shift from C-corporations to S-corporations in response to TRA86, it has been suggested that S-corporation income should be placed with capital rather than business income.

⁷⁸Saez and Veall (2007) find that in Canada top wages increased similarly to the USA without the same changes in fiscal policy. They argue that due to the threat of migration to the US, Canadian salaries had to keep up with US top wages, which thus experienced real increases.

in the long-run (Atkinson et al., 2011).

Concerning the last tier of the Slemrod hierarchy, the real effects of the tax cuts on economic activity and thus top shares have been limited (Saez et al., 2012).⁷⁹ Figure 2.2 shows, and this is confirmed for smaller top fractiles (Piketty and Saez, 2007b), that top shares started to rise before TRA86 and accelerated in the late 1990s despite the rise in top tax rates in 1993 (Goolsbee, 2000). Recently, Yagan (2013) finds no real effects of the substantial 2003 dividend tax cut in the US using a quasi-experimental design.

2.A.4 Comparing market income with Haig-Simons income

I measure income reported on tax returns, which has been criticised for not conforming to the preferred Haig-Simons definition of income. According to Haig-Simons, income during a year is defined as the net increase in one’s ability to consume.⁸⁰ Note that this income definition is different from disposable income, which excludes capital gains altogether (Canberra Group, 2011)

Two particular aspects of my income definition have been criticised. First, I exclude any income that is not taxable. I thus exclude public welfare payments, tax-exempt interest⁸¹, in-kind income, imputed rents, non-taxable fringe benefits such as health insurance, gifts, inheritances, and distributions from trusts and life insurance policies. CBO (2012) impute non-taxable income (e.g. transfers and in-kind income) from the CPS into the IRS data. Reassuringly for my methodology, market income accounts for 80% of CBO-incomes, and an even larger proportion at the top of the distribution, which means that the top shares series are quite similar (Bivens and Mishel, 2013).

⁷⁹Auerbach and Slemrod (1997) conclude that “it is likely that TRA86 is the leading cause of the surge in the reported income of the affluent” [p. 616]. The availability of newer data showing the strong increase in top shares in the late 1990s might explain why Slemrod changed his opinion between his 2012 (Saez et al., 2012) and his 1997 (Auerbach and Slemrod, 1997) papers.

⁸⁰Haig (1921) defined income as “the *money value of the net accretion to one’s economic power between two points of time*” [emphasis in original, p. 7]. Simons (1938) describes income as “the algebraic sum of (1) the market value of rights exercised in consumption and (2) the change in the value of the store of property rights between the beginning and end of the period in question” [p. 50]. Hill and Hill (2003) define it “as the sum of consumption plus the actual change in wealth over the period” [p. 611].

⁸¹The interest received on state and local bonds is not taxable. This income is still reportable on a tax return but only after 1987.

Furthermore, I only include taxable capital gains, which exclude the gains on the sale of a primary residence (subject to an upper limit) or gains on inherited property under a stepped-up basis.⁸² However, the importance of these income sources as a share of income declines with income; e.g. capital gains on primary housing (Armour et al., 2013).

The second critique is that my income definition recognises capital gains upon realisation rather than accrual, although the Haig-Simons definition clearly includes accrued (yet unrealised) capital gains. On the other hand, the accruals approach is difficult to implement in thin markets, or might result in temporary capital gains due to stock market booms. Furthermore, it might simply be wrong, because these gains cannot be used to finance consumption, precisely because they are temporary.

In Norway, Aaberge and Atkinson (2010) find that top shares using a more comprehensive Hicksian measure of capital income show a slower increase than the standard series based on realisations.⁸³ I cannot follow their approach because US tax data do not record asset holdings. In Sweden, Roine and Waldenström (2012) account for the short-run fluctuations in capital gain realisations by taking long-run averages of capital gains. However, this requires panel data which are not available for the US.

The main source of information about household wealth in the US is the Survey of Consumer Finances (SCF) by the Federal Reserve Board. Armour et al. (2013) use CPS data and impute capital gains from IRS and SCF, which relies on a number of strong assumptions.⁸⁴ Including non-taxable capital gains and transfers, and after adjusting

⁸²Under a stepped-up provision, the fair market value at the time of death becomes the basis for assessing any future capital gains. Hence any capital gains by the deceased at the time of death are never taxed.

⁸³Hicksian capital income uses data on wealth holdings in combination with long-run data on returns for different asset classes. Hill and Hill (2003) define Hicksian income as the “maximum amount that can be consumed while maintaining wealth intact” [p. 607], which includes all capital gains (under perfect foresight). It is very similar to a Haig-Simons definition of income, with the exception of unexpected capital gains which are excluded from generalised-Hicksian income, while Haig-Simons includes them.

⁸⁴They construct tax units and taxable income in the CPS. To each percentile of taxable income in the CPS they assign the mean capital gains observed in the IRS data for that percentile. They thus ignore any inequality in capital gains within the percentile of taxable income. They also need to assume that the relationship between taxable income and capital gains is the same in the CPS and IRS datasets. Burkhauser et al. (2012) show that the CPS understates top incomes, so the percentiles of taxable income may not be well aligned.

for household size and taxes, they claim that the faster increase in incomes at the top disappears. However, these results have been criticised (Edsall, 2013; Thompson and Smeeding, 2013).⁸⁵ Most importantly, their treatment of capital gains at accrual, both negative and positive, is misleading for a number of reasons: First, they compare changes between 1989 and 2007, both of which are peak years. Second, with the SCF data, they assign everyone the same rate of return, which might be misleading if the rich have access to better investment advisors. Third, this rate of return comes from the Dow-Jones, which makes the capital gains volatile. Fourth, and most importantly, they use the full value of accrued capital gains whether positive or negative. For example, consider a tax unit who receives an income of \$200k and owns a house worth \$4M. During 2007, when house prices fell by almost 5% (Edsall, 2013), the tax unit would have a \$200k capital loss, and consequently zero income, which is unlikely to reflect their true living standard.

Other studies based on the SCF and using a more complete measure of capital incomes support the results from the IRS data. Wolff and Zacharias (2009) compute capital incomes including for example the lifetime annuity value of non-home wealth and imputed rents for owner-occupiers. Their top shares are higher than Piketty and Saez (2007b) but follow a similar trend over time. This is also confirmed by Smeeding and Thompson (2011).

2.A.5 Using before-tax incomes

Most of the literature on top income shares in the US focusses on gross incomes.⁸⁶ This appears surprising given that (1) other studies of income inequality typically concentrate on disposable income, and (2) the tax micro data contain information on the taxes paid. The literature on historic top income shares uses tabulations of tax units by before-tax income. These tabulations might also contain total taxes paid for

⁸⁵First, as mentioned above, the CPS is not suitable for studying top incomes. Second, imputing the value of government transfers which are most important at the bottom of the distribution, might be misleading, particularly for health care which is very costly in the US (Baker, 2013).

⁸⁶In the UK, Atkinson (2007a) finds that the U-shaped pattern of top shares is robust to an after-tax analysis, and that the recent rise in top shares becomes even more stark when taxes are considered.

the same before-tax income ranges. Using this in an after-tax analysis would, however, ignore the re-ranking effects of income taxation which can be substantial (Atkinson, 2007a).

For the purposes of this chapter, it is important to realise that the information in the tax files only refers to federal income taxes. In 2011, 46% of tax units paid no federal income taxes, for example because of taxpayer and dependent exemptions, the Earned Income Tax Credit, or deductions such as for education expenses (Johnson et al., 2011). Of course, these tax units might still pay other important taxes, such as federal payroll taxes, state income taxes and local property taxes. To capture the redistributive effects of the *federal* tax system more fully, Piketty and Saez (2007a) use a model to impute a tax liability. CBO (2012) extends their approach using a more comprehensive income measure which imputes non-taxable income from the CPS. However, both studies ignore state⁸⁷ and local taxes, and abstract from any re-ranking effects induced by the tax system.⁸⁸ In this chapter, I follow the earlier literature and use gross income, because (a) I want to avoid any additional assumptions necessary for the imputation, and (b) it is not clear how to allocate taxes to labour and capital in the second part of the chapter.

⁸⁷These taxes account for around one third of US taxes and are somewhat regressive (Piketty and Saez, 2007a). State taxes cannot be imputed into the IRS data, because the state of residence is removed at least for high income tax units. Furthermore, while state taxes are deductible from federal income taxes, they would only be listed separately if the tax unit chooses to itemise deductions.

⁸⁸That is, both studies analyse the distribution of pre- and post-tax income for tax units ranked according to pre-tax income (typically excluding capital gains). However, the tax code could also lead to a re-ranking of tax units so the top 1% of gross incomes would be different from the top 1% of after-tax incomes. Tax units with the same pre-tax income pay different taxes for example because of differences in tax unit size and composition, or income composition, or because some simply have a better tax advisor (Atkinson, 1980).

2.B Appendix Tables

Table 2.B.1: Total income share excl. capital gains (in %): Comparing to Saez (2013)

year	P90-95		P95-99		P99-99.5		P99.5-99.9		P99.9-99.99		P99.99-100		avg. diff.
	own	Saez	own	Saez	own	Saez	own	Saez	own	Saez	own	Saez	
1960	10.87	11.15	12.50	12.15	2.74	2.84	3.40	3.42	1.54	1.50	0.59	0.60	
1962	11.10	11.10	12.77	12.67	2.82	2.87	3.47	3.42	1.53	1.42	0.57	0.56	
1964	11.12	11.02	12.75	12.60	2.79	2.69	3.40	3.36	1.48	1.44	0.56	0.53	
1966	11.00	10.99	12.67	12.62	2.79	2.78	3.42	3.43	1.54	1.55	0.60	0.60	0.01
1967	11.00	10.97	12.70	12.65	2.80	2.80	3.47	3.47	1.56	1.56	0.59	0.60	0.01
1968	11.00	11.01	12.65	12.62	2.78	2.77	3.43	3.43	1.55	1.56	0.58	0.58	0.00
1969	11.02	11.14	12.56	12.66	2.70	2.71	3.27	3.30	1.43	1.45	0.54	0.55	-0.05
1970	11.14	11.13	12.64	12.58	2.66	2.65	3.22	3.22	1.40	1.41	0.52	0.53	0.01
1971	11.30	11.26	12.78	12.71	2.68	2.66	3.22	3.21	1.39	1.40	0.51	0.52	0.02
1972	11.28	11.25	12.69	12.62	2.68	2.66	3.19	3.18	1.39	1.40	0.52	0.52	0.02
1973	11.29	11.28	12.89	12.83	2.69	2.67	3.19	3.18	1.38	1.39	0.49	0.50	0.01
1974	11.29	11.32	12.93	12.91	2.73	2.71	3.30	3.30	1.53	1.54	0.55	0.56	0.00
1975	11.63	11.60	13.11	13.02	2.73	2.70	3.30	3.27	1.47	1.48	0.55	0.56	0.03
1976	11.59	11.57	13.05	12.96	2.69	2.66	3.23	3.21	1.45	1.46	0.56	0.56	0.02
1977	11.60	11.60	13.04	12.93	2.69	2.65	3.24	3.21	1.47	1.48	0.56	0.57	0.03
1978	11.58	11.58	13.02	12.91	2.69	2.65	3.25	3.22	1.49	1.50	0.57	0.58	0.03
1979	11.52	11.52	12.92	12.80	2.69	2.65	3.26	3.23	1.54	1.54	0.61	0.62	0.03
1980	11.70	11.70	13.12	12.99	2.71	2.67	3.31	3.28	1.57	1.58	0.65	0.65	0.03
1981	11.81	11.75	13.13	12.94	2.65	2.60	3.21	3.20	1.54	1.57	0.64	0.66	0.04
1982	11.99	11.82	13.32	13.01	2.72	2.66	3.34	3.28	1.68	1.68	0.77	0.77	0.10
1983	12.12	11.91	13.55	13.19	2.72	2.66	3.39	3.33	1.74	1.74	0.86	0.87	0.11
1984	12.04	11.85	13.52	13.21	2.72	2.67	3.44	3.39	1.84	1.85	0.97	0.98	0.10
1985	12.06	11.87	13.61	13.28	2.76	2.70	3.50	3.48	1.98	1.94	0.96	0.97	0.10
1986	12.20	11.98	13.81	13.46	2.81	2.75	3.55	3.51	1.87	1.87	0.99	1.00	0.11
1987	11.99	11.99	13.85	13.74	3.04	2.98	4.10	4.04	2.44	2.43	1.30	1.30	0.04
1988	11.69	11.68	13.89	13.78	3.26	3.20	4.80	4.75	3.23	3.22	1.99	1.99	0.04
1989	11.82	11.81	14.16	14.05	3.29	3.24	4.69	4.63	3.02	3.00	1.75	1.74	0.04
1990	11.80	11.78	14.20	14.07	3.32	3.27	4.88	4.82	3.09	3.07	1.83	1.83	0.05
1991	12.03	11.95	14.52	14.26	3.36	3.27	4.68	4.54	2.83	2.75	1.61	1.61	0.11
1992	11.99	11.94	14.63	14.40	3.46	3.37	4.99	4.90	3.22	3.20	2.03	2.02	0.08
1993	12.12	12.07	14.81	14.59	3.45	3.37	4.82	4.74	3.00	2.98	1.74	1.74	0.08
1994	12.08	12.09	14.81	14.65	3.47	3.40	4.82	4.74	3.00	2.97	1.74	1.73	0.05
1995	12.09	12.08	14.94	14.93	3.54	3.54	5.01	5.01	3.17	3.17	1.82	1.82	0.00
1996	12.00	12.00	15.06	15.05	3.62	3.62	5.16	5.16	3.35	3.35	1.97	1.97	0.00
1997	11.87	11.87	15.08	15.08	3.65	3.65	5.32	5.32	3.61	3.61	2.19	2.20	0.00
1998	11.76	11.77	15.06	15.06	3.68	3.68	5.41	5.41	3.80	3.79	2.41	2.41	0.00
1999	11.74	11.70	15.13	15.10	3.73	3.72	5.54	5.52	4.02	4.00	2.64	2.63	0.02
2000	11.68	11.60	15.19	15.02	3.76	3.71	5.70	5.66	4.29	4.29	2.84	2.84	0.06
2001	11.95	11.83	15.25	15.03	3.73	3.66	5.52	5.45	3.88	3.86	2.41	2.40	0.08
2002	12.07	12.00	15.36	15.37	3.70	3.73	5.39	5.33	3.64	3.63	2.19	2.30	
2003	12.17	12.11	15.38	15.44	3.72	3.75	5.40	5.36	3.68	3.67	2.35	2.44	
2004	12.04	11.94	15.44	15.37	3.80	3.83	5.68	5.60	4.06	4.04	2.72	2.87	
2005	12.03	11.82	15.64	15.44	3.98	3.96	6.09	5.96	4.57	4.47	3.07	3.29	
avg. diff.	0.04		0.13		0.04		0.03		0.00		0.00		
min. diff.	-0.12		-0.10		-0.01		-0.03		-0.03		-0.02		
max. diff.	0.22		0.36		0.10		0.13		0.09		0.01		

Using full sample of tax records. Following Piketty and Saez (2007b) methodology, but ranking by gross market income, not AGI. Column “Saez” refers to the most recent figures from Saez (2013).

Differences (average, maximum, minimum) are between the respective columns (in percentage points).

Differences only calculated for years in which Saez (2013) uses micro data.

Table 2.B.2: Total income share incl. capital gains (in %): Comparing to Saez (2013)

year	P90-95		P95-99		P99-99.5		P99.5-99.9		P99.9-99.99		P99.99-100		avg. diff.
	own	Saez	own	Saez	own	Saez	own	Saez	own	Saez	own	Saez	
1960	10.62	10.90	12.57	12.54	2.88	2.91	3.71	3.88	2.04	2.07	1.28	1.17	
1962	10.92	10.90	12.78	12.86	2.89	2.89	3.72	3.86	1.95	2.03	1.14	1.16	
1964	10.87	10.92	12.87	13.02	2.96	3.09	3.79	4.01	2.02	2.07	1.23	1.30	
1966	10.76	10.75	12.77	12.74	2.96	2.96	3.81	3.83	2.08	2.10	1.28	1.29	0.00
1967	10.76	10.74	13.00	12.96	3.05	3.06	3.98	4.00	2.24	2.26	1.40	1.42	0.00
1968	10.69	10.70	12.96	12.94	3.08	3.07	4.10	4.12	2.38	2.41	1.59	1.61	-0.01
1969	10.73	10.85	12.64	12.73	2.89	2.91	3.71	3.75	2.10	2.14	1.52	1.56	-0.06
1970	10.97	10.96	12.68	12.64	2.78	2.77	3.47	3.48	1.76	1.78	0.98	1.00	0.00
1971	11.10	11.08	12.89	12.86	2.84	2.84	3.56	3.58	1.85	1.87	1.10	1.11	0.00
1972	11.03	11.07	12.88	12.88	2.85	2.86	3.63	3.65	1.92	1.95	1.16	1.18	-0.02
1973	11.09	11.12	13.08	13.05	2.85	2.86	3.53	3.54	1.80	1.82	0.93	0.94	-0.01
1974	11.15	11.19	12.99	13.00	2.82	2.82	3.56	3.58	1.82	1.85	0.86	0.88	-0.02
1975	11.47	11.45	13.20	13.11	2.83	2.80	3.52	3.50	1.71	1.72	0.84	0.85	0.02
1976	11.45	11.44	13.18	13.11	2.82	2.79	3.48	3.47	1.72	1.73	0.85	0.86	0.01
1977	11.44	11.46	13.19	13.10	2.83	2.81	3.52	3.51	1.77	1.78	0.91	0.92	0.01
1978	11.45	11.45	13.20	13.09	2.82	2.79	3.53	3.51	1.78	1.79	0.84	0.86	0.02
1979	11.30	11.28	13.11	12.97	2.89	2.85	3.69	3.67	2.06	2.07	1.36	1.37	0.03
1980	11.45	11.47	13.25	13.15	2.90	2.87	3.75	3.74	2.12	2.13	1.26	1.28	0.01
1981	11.55	11.51	13.20	13.02	2.82	2.78	3.65	3.67	2.13	2.20	1.31	1.37	0.02
1982	11.64	11.50	13.30	13.04	2.87	2.83	3.81	3.79	2.43	2.44	1.71	1.73	0.07
1983	11.71	11.53	13.57	13.30	2.95	2.92	4.02	4.01	2.70	2.74	1.86	1.88	0.07
1984	11.61	11.45	13.53	13.30	2.96	2.95	4.05	4.06	2.80	2.83	2.13	2.15	0.06
1985	11.59	11.44	13.68	13.45	3.07	3.04	4.29	4.31	3.04	3.08	2.21	2.24	0.05
1986	11.36	11.14	13.85	13.57	3.31	3.30	4.84	5.22	3.75	4.05	3.32	3.34	-0.03
1987	11.69	11.71	13.96	13.88	3.26	3.21	4.59	4.55	3.01	2.99	1.90	1.91	0.03
1988	11.34	11.34	13.89	13.80	3.45	3.40	5.33	5.29	3.94	3.94	2.86	2.86	0.03
1989	11.54	11.54	14.15	14.06	3.45	3.41	5.13	5.08	3.56	3.54	2.46	2.45	0.03
1990	11.58	11.57	14.19	14.08	3.44	3.39	5.17	5.12	3.50	3.49	2.34	2.33	0.04
1991	11.86	11.82	14.53	14.36	3.45	3.38	4.93	4.86	3.17	3.17	1.96	1.96	0.06
1992	11.80	11.76	14.61	14.39	3.55	3.47	5.24	5.17	3.59	3.57	2.47	2.46	0.07
1993	11.88	11.85	14.79	14.60	3.54	3.46	5.12	5.05	3.43	3.41	2.32	2.32	0.07
1994	11.86	11.89	14.79	14.66	3.56	3.50	5.10	5.03	3.43	3.41	2.31	2.29	0.04
1995	11.83	11.89	14.91	14.99	3.64	3.66	5.34	5.36	3.72	3.74	2.45	2.46	-0.04
1996	11.64	11.73	14.96	15.07	3.75	3.78	5.62	5.67	4.15	4.18	3.04	3.06	-0.05
1997	11.41	11.50	15.00	15.13	3.83	3.86	5.93	5.97	4.62	4.66	3.50	3.53	-0.06
1998	11.23	11.30	14.91	15.01	3.88	3.91	6.14	6.18	5.04	5.07	3.90	3.92	-0.05
1999	11.18	11.25	15.07	15.17	3.98	4.01	6.37	6.41	5.37	5.41	4.19	4.21	-0.05
2000	11.04	11.00	14.98	15.09	3.98	4.07	6.57	6.58	5.79	5.80	4.81	5.07	-0.07
2001	11.57	11.47	15.12	15.13	3.84	3.90	5.99	5.96	4.68	4.67	3.51	3.70	-0.02
2002	11.78	11.75	15.28	15.20	3.81	3.83	5.76	5.70	4.22	4.20	2.99	3.14	
2003	11.82	11.76	15.29	15.24	3.86	3.86	5.85	5.80	4.39	4.38	3.28	3.49	
2004	11.52	11.45	15.28	15.20	4.00	4.01	6.34	6.28	5.14	5.12	4.09	4.34	
2005	11.38	11.18	15.47	15.24	4.22	4.15	6.90	6.79	5.98	5.86	4.68	5.13	
avg. diff.	0.02		0.07		0.01		-0.01		-0.02		-0.03		
min. diff.	-0.11		-0.12		-0.08		-0.38		-0.31		-0.26		
max. diff.	0.22		0.28		0.08		0.08		0.03		0.01		

Using full sample of tax records. Following Piketty and Saez (2007b) methodology, but ranking by gross market income, not AGI. Column “Saez” refers to the most recent figures from Saez (2013).

Differences (average, maximum, minimum) are between the respective columns (in percentage points).

Differences only calculated for years in which Saez (2013) uses micro data.

Table 2.B.3: Total income share incl. capital gains (in %): Comparing Part II sample

year	P90-95		P95-99		P99-99.5		P99.5-99.9		P99.9-100		avg. diff.
	old	new	old	new	old	new	old	new	old	new	
1960	10.62	10.68	12.57	11.69	2.88	2.44	3.71	3.22	3.32	3.77	0.26
1962	10.93	10.29	12.53	11.94	2.78	2.67	3.57	3.42	2.71	2.61	0.32
1964	10.91	10.35	12.67	12.07	2.86	2.74	3.66	3.47	2.85	2.75	0.31
1966	10.82	10.44	12.64	12.25	2.87	2.80	3.66	3.54	2.89	2.83	0.20
1967	10.82	10.48	12.87	12.54	2.96	2.90	3.82	3.73	3.09	3.03	0.17
1968	10.79	10.53	12.86	12.61	2.99	2.94	3.89	3.82	3.24	3.21	0.13
1969	10.84	10.65	12.45	12.25	2.77	2.73	3.54	3.49	2.93	2.93	0.10
1970	11.01	10.64	12.45	12.08	2.66	2.59	3.30	3.21	2.32	2.27	0.19
1971	11.16	10.71	12.72	12.23	2.72	2.63	3.36	3.25	2.42	2.36	0.24
1972	11.14	10.76	12.67	12.29	2.71	2.63	3.38	3.30	2.48	2.44	0.19
1973	11.13	10.83	12.79	12.49	2.72	2.66	3.31	3.23	2.27	2.24	0.15
1974	11.15	10.92	12.69	12.45	2.68	2.62	3.31	3.25	2.30	2.27	0.12
1975	11.52	11.10	12.88	12.42	2.66	2.57	3.27	3.16	2.16	2.11	0.23
1976	11.50	11.11	12.91	12.50	2.65	2.57	3.21	3.13	2.17	2.13	0.20
1977	11.51	11.12	12.89	12.48	2.65	2.58	3.23	3.15	2.23	2.20	0.20
1978	11.55	11.21	12.90	12.55	2.64	2.58	3.21	3.14	2.19	2.16	0.17
1979	11.42	11.13	12.79	12.49	2.66	2.62	3.30	3.25	2.69	2.69	0.14
1980	11.57	11.24	12.90	12.54	2.67	2.60	3.30	3.24	2.58	2.56	0.16
1981	11.66	11.32	12.80	12.42	2.56	2.50	3.13	3.07	2.50	2.48	0.17
1982	11.79	11.34	12.90	12.43	2.57	2.51	3.22	3.15	2.99	2.96	0.22
1983	11.86	11.37	13.13	12.61	2.61	2.53	3.32	3.25	3.31	3.27	0.24
1984	11.75	11.35	13.07	12.63	2.59	2.54	3.33	3.27	3.48	3.46	0.20
1985	11.71	11.34	13.17	12.79	2.64	2.61	3.56	3.49	3.89	3.86	0.18
1986	11.64	11.26	13.40	13.00	2.90	2.86	3.97	3.90	4.89	4.86	0.19
1987	11.87	11.57	13.48	13.16	2.86	2.81	3.78	3.72	3.64	3.60	0.15
1988	11.59	11.37	13.46	13.22	3.03	2.98	4.40	4.34	5.36	5.21	0.14
1989	11.69	11.48	13.75	13.51	3.09	3.05	4.43	4.37	4.80	4.71	0.13
1990	11.70	11.47	13.79	13.52	3.11	3.07	4.52	4.46	4.83	4.73	0.14
1991	11.87	11.62	14.14	13.86	3.23	3.17	4.50	4.43	4.42	4.31	0.15
1992	11.85	11.49	14.27	13.86	3.31	3.25	4.83	4.74	5.31	5.17	0.21
1993	11.94	11.55	14.47	14.05	3.33	3.26	4.75	4.65	5.02	4.91	0.22
1994	11.92	11.55	14.47	14.06	3.37	3.29	4.77	4.67	5.07	4.96	0.21
1995	11.85	11.53	14.66	14.31	3.51	3.45	5.04	4.95	5.55	5.43	0.19
1996	11.71	11.43	14.81	14.51	3.63	3.58	5.36	5.27	6.46	6.36	0.16
1997	11.53	11.26	14.88	14.61	3.71	3.67	5.63	5.56	7.26	7.14	0.15
1998	11.35	11.13	14.83	14.61	3.77	3.73	5.86	5.79	7.88	7.77	0.13
1999	11.35	11.15	15.04	14.83	3.90	3.87	6.10	6.04	8.42	8.31	0.12
2000	11.19	11.04	14.98	14.80	3.94	3.91	6.34	6.29	9.26	9.13	0.11
2001	11.57	11.30	14.89	14.60	3.76	3.71	5.77	5.74	7.32	7.34	0.12
2002	11.71	11.38	14.89	14.47	3.68	3.62	5.54	5.47	6.55	6.54	0.18
2003	11.82	11.42	14.93	14.48	3.70	3.63	5.61	5.52	6.91	6.91	0.20
2004	11.61	11.25	15.15	14.74	3.91	3.85	6.11	6.01	7.94	7.79	0.22
2005	11.44	11.10	15.29	14.93	4.10	4.05	6.57	6.49	9.22	9.05	0.20
avg. diff.	0.33		0.37		0.07		0.09		0.05		
min. diff.	-0.06		0.18		0.03		0.03		-0.45		
max. diff.	0.64		0.88		0.43		0.49		0.17		

Using Part II sample (i.e. excl. observations with negative labour, capital, or business).

Columns “old” adjust for non-filers and use total income (as Table 2.B.2), but use the new sample. Columns “new” do not adjust for non-filers and use sum of component incomes.

Differences (average, maximum, minimum) are between the respective columns (in pp).

Chapter 3

Enterprise Size and Earnings: Evidence from Urban Ghana

Introduction

This chapter attempts to explain the observed heterogeneity in labour market earnings in urban Ghana. In particular, I examine the enterprise size premium which has been observed in a number of cross-sectional datasets. In Ghana I find that without controls for individual fixed effects, an individual working in a firm of size 100 earns approximately 64% more than an individual working in a firm of size 5.¹ The size premium in self-employment which has received much less attention in the literature is approximately 50% larger than the premium for wage-earners. I investigate the relationship between earnings and enterprise size in a common framework across the wage- and self-employed.² Size can explain a significant fraction of the earnings heterogeneity within each sector. In a labour market like Ghana, wage-employment especially in

¹This differential is substantially larger than the differential found in the US (Brown and Medoff, 1989). This conclusion also holds in other studies on developing countries (Funkhouser, 1998; Velenchik, 1997).

²For the self-employed, the size variable refers to the number of employees and the dependent variable is the income generated in self-employment. Sandefur et al. (2007) and Falco et al. (2011) investigate the enterprise size effect in an earnings regression pooling the wage- and self-employed. At the time of writing my MPhil thesis I was not aware of these contributions, which are quite different in any case. Sandefur et al. (2007) only exploit the cross-sectional variation. There are a number of difference with Falco et al. (2011), as (1) I cover a longer time period; (2) I estimate different earnings regressions for the wage- and self-employed; (3) I use a Difference-GMM estimator; and (4) I check for robustness using the matched firm-worker data.

large firms is rare and a large fraction of the labour force is self-employed.³ Of course, the mechanisms which connect scale to earnings are quite different across the two types of employment. For example, earnings in self-employment would also include returns to physical capital, in addition to human capital (Sandefur et al., 2007).

I investigate the firm size effect in urban Ghana using two quite different panel datasets. The first is a matched firm-worker dataset which oversamples large firms. The second is a representative labour force survey which is more informative about small and medium firms. I find evidence of a significant firm size premium in both datasets even after controlling for time-invariant unobservables. These unobservables account for between one half and two thirds of the cross-sectional premium. This is also robust to allowing for endogeneity by using Difference- and System-GMM estimators with some specifications of the instrument set. The labour force survey provides data on the self-employed which allow me to study the size effect in self-employment. In this sector, the cross-section size premium seems to be driven by time-invariant unobservables, which is robust to allowing for endogeneity by using dynamic panel estimators.

This chapter adds to the literature in three ways. First, I control for time-invariant characteristics using two panel datasets. Second, I allow for endogeneity in these estimates by using Difference- and System-GMM estimators. Third, I estimate the size effect in wage- and self-employment separately, given that the mechanisms are quite different across these sectors. This chapter is structured in four main parts. The first section presents the analytical framework, reviews the existing literature and describes the panel estimators. Section 3.2 introduces the two datasets and provides summary statistics. In Section 3.3 I present the results. Section 3.4 concludes.

³In 1987, 46% (28%) of manufacturing employees in Ghana were employed in firms with less than 100 (20) workers, compared with 18% (6%) for the US (Oi and Idson, 1999; Sandefur et al., 2007). In Ghana, the share of employment in small firms grew to 67% in firms with less than 100 and 48% in firms with less than 20 workers by 2003. Figure 3.A.1 shows a robust positive relationship between log GDP per capita and firm size across countries. At the same time, non-agricultural self-employment is the fastest growing part of the urban labour market (Rankin et al., 2010). For Ghana, the share of self-employment in the non-agricultural labour force increased from 50% in 1988 to 63% in 1999 (Kingdon et al., 2006), compared with an average of 7% in the US during this period (Blanchflower, 2000). Figure 3.A.2 shows a significant positive relationship between the logarithm of GDP and the percentage of male workers who receive a wage or salary.

3.1 Econometric framework

3.1.1 Literature review

In a perfectly competitive labour market without frictions, there should not be any premium to working in a larger firm; otherwise production should take place in a single firm. Therefore, in this type of model the significant size premium is explained by inadequate controls for worker skills or differences in the non-pecuniary benefits of working in large or small firms, which might both be unobserved. However, a real size premium can arise in more elaborate labour market models. In a model with search frictions (e.g. incomplete information about the wages on offer) and heterogeneous employers, highly productive employers contact more workers and offer them higher wages, thereby ending up with more workers (Mortensen, 2003).⁴ Matching or sorting models allow both heterogeneous employers and workers (Shimer and Smith, 2000; Caselli and Coleman, 2006), and under certain conditions assortative matching can arise, i.e. high-productivity firms match with high ability individuals. For example, if larger firms have better hiring practices they could be better at spotting unobserved worker characteristics (Weiss and Landau, 1984).

The remainder of this section presents the econometric framework and reviews the existing estimates of the firm size effect. The firm size effect can be estimated consistently using matched employer-employee data. Self-employed individuals can be thought of as being matched with their own enterprise-specific fixed effect, so their size effect can be estimated consistently from labour force data. In other words, the interpretation of the fixed effect is very different in the wage- and self-employed samples.⁵ The size effect is typically estimated in a standard Mincerian earnings function (Mincer,

⁴Increasing marginal recruitment costs are essential for this equilibrium to emerge. Mortensen (2003) is silent about the sources of this heterogeneity. They could arise from differences in managerial skill, fixed factors of production or the costs of capital. If monitoring and other enforcement mechanisms are more difficult in larger firms, these firms are more likely to pay efficiency wages, i.e. the profit-maximising wage exceeds the market wage (Bulow and Summers, 1986; Ringuedé, 1998). Using Ghanaian data, Söderbom and Teal (2004), for example, find that larger firms have lower costs of capital. Also see the discussion in the conclusion to this chapter.

⁵This is one reason why I estimate the size effects in wage- and self-employment separately, in contrast to Sandefur et al. (2007) and Falco et al. (2011).

1974)

$$y_{it} = x_{it}\beta + s_{it}\delta + \kappa_i + \varepsilon_{it} \quad (3.1)$$

where y_{it} is the logarithm of earnings of individual $i = 1, \dots, N$ at time $t = 1, \dots, T$; x_{it} is a $(1 \times P)$ vector of time-varying personal and enterprise characteristics, such as tenure, age or the firm's capital stock; β is the $(P \times 1)$ vector of corresponding coefficients; κ_i is the time-invariant individual fixed effect; ε_{it} is a mean-zero error term which is orthogonal to the included regressors. I am interested in estimating δ , the coefficient on s_{it} , the natural logarithm of enterprise size. Enterprise size is defined as firm size in wage-employment and the number of employees in self-employment. Since both y_{it} and s_{it} are measured in natural logarithms, δ denotes the elasticity of earnings with respect to enterprise size.

The κ_i term incorporates both person- and enterprise-specific time-invariant characteristics. These might be observable, such as education or gender, or might be unobservable, such as ability.⁶ κ_i can be decomposed as $\kappa_i = \theta_i + \psi_i$, where ψ_i and θ_i are respectively the enterprise- and person-specific time-invariant characteristics, referred to as the “pure person” and “pure firm” effects by Abowd et al. (1999). In both the matched firm-worker data and the data for the self-employed in the labour force survey, I can control for both of these effects and thus estimate δ consistently. However, because I do not observe movement across firms by the same individual, I cannot disentangle the person and firm effects (in contrast to Abowd et al., 1999). Furthermore, this implies that the estimated size coefficient is likely to be a conservative estimate compared with the general population.⁷ Consistent estimation of Equation 3.1 requires that the pure error term ε_{it} is orthogonal to the included regressors, i.e. there is no correlation between time-varying unobservables and the regressors of interest. This assumes that the choice between wage- and self-employment depends only on time-

⁶ Education refers to formal schooling, which I assume is unchanged once the individual has entered the labour force.

⁷ Brown and Medoff (1989) and Winter-Ebmer and Zweimüller (1999) estimate a smaller size effect for “stayers” than for “movers”.

invariant characteristics.⁸

Using the same matched employer-employee data, Söderbom et al. (2005) find a significant size effect for the wage-employed even after controlling for person and firm heterogeneity. I extend their analysis of the wage-employed by using a longer panel and allowing for a more general econometric specification. Using French data, Abowd et al. (1999) find that the individual fixed effect explains approximately 75% of the effect of firm size on wages, while firm effects explain little. In the US, Abowd and Kramarz (2000) find the opposite result that firm effects are quantitatively more important. However, even with controls for unobservables, there remains a residual size-wage effect in both papers.⁹

The literature prior to Abowd et al. (1999) used cross-section data and worker panel data, and thus did not estimate Equation 3.1 consistently. Early papers on the firm size effect in cross-section data estimated the following equation, which omits the individual fixed effect

$$y_{it} = x_{it}\beta^* + s_{it}\delta^* + u_{it} \quad (3.2)$$

where $u_{it} = \kappa_i + \varepsilon_{it}$ is the error term. As mentioned in the introduction, these studies found a strongly significant effect of a similar magnitude as the gender pay gap (e.g. Oi and Idson, 1999). The estimated firm size coefficient is given by (as derived in Abowd et al., 1999)¹⁰

⁸This implies “exogenous switching” (Maddala, 1983), i.e. the error terms in the outcome and selection equations are uncorrelated. The same applies to selection into employment. The selection between wage- and self-employment is particularly important in this context given the large share of self-employment.

⁹This contrasts with the industry premium which disappears almost completely after controlling for individual and firm heterogeneity in the French data, as corroborated by Goux and Maurin (1999).

¹⁰As usual, an estimator is denoted by a “hat”. A single star refers to estimators which omit both person and firm effects, while a double star refers to the case where only the firm effect is omitted (see Equation 3.4). Stacking over t and i , I can write Equation 3.2 as $y = X\beta^* + S\delta^* + u$. Defining matrices W and ζ^* as $W = \{XS\}$ and $\zeta^* = \begin{Bmatrix} \beta^* \\ \delta^* \end{Bmatrix}$, this can be rewritten as $y = W\zeta^* + u$. The true data generating process is $y = X\beta + S\delta + G\kappa + \varepsilon$, where G is the $NT \times N$ matrix of indicators for individuals $i = 1, \dots, N$. Using standard results for OLS estimators, I can write the expression for the estimator of ζ^* as $(W'W)\hat{\zeta}^* = W'y$. Substituting for W and ζ^* , and multiplying out the matrices, I get the two standard normal equations. Solving these two equations simultaneously and defining $M_X = [I - X(X'X)^{-1}X']$ as the standard annihilator matrix (Hayashi, 2000), yields $\hat{\delta}^* = (S'M_X S)^{-1}S'M_X y$. The bias from omitting the individual fixed effect κ can be derived by

$$\hat{\delta}^* = \delta + (S' M_X S)^{-1} S' M_X G \kappa \quad (3.3)$$

where $M_X = I - X(X'X)^{-1}X'$ and G is the $(NT \times N)$ matrix of indicators for N individuals. The OLS estimator of δ^* is inconsistent unless all the included regressors are orthogonal to the omitted individual fixed effect.¹¹ I find that the estimated size effect in the cross-section is substantially larger than using a fixed effects or first-differences estimator, which holds in both wage- and self-employment.

A number of studies, such as Brown and Medoff (1989), use worker panel data to control for individual unobservables, θ_i , but cannot control for firm effects. The same issue arises when I estimate the firm size effect using the wage-employees in the labour force data. With worker panel data the following equation is estimated

$$y_{it} = x_{it}\beta^{**} + s_{it}\delta^{**} + \theta_i + v_{it} \quad (3.4)$$

where $v_{it} = \psi_i + \varepsilon_{it}$. Similar to above, the estimated size coefficient can be written as¹²

$$\hat{\delta}^{**} = \delta + (S' M_{[X \ G]} S)^{-1} S' M_{[X \ G]} G \psi$$

For $\hat{\delta}^{**}$ to be a consistent estimator of δ , firm size would have to be orthogonal to $G\psi$ *conditional* on the time-varying personal characteristics, X , and pure person effects, G , as represented by the presence of the $M_{[X \ G]}$ term.¹³ This rules out the possibility that, conditional on pure person effects, more productive firms (captured by firm unobservables) are larger. In my application, the estimated firm size effect in the labour force data is very similar to that found in the matched data. However, the results might not be comparable because (1) the matched data in contrast to the labour

substituting the true data generating process for y , by recognising that $M_X X = 0$ and that the error term ε is orthogonal to each of the regressors.

¹¹The second term in Equation 3.3 captures the partial correlations between the included variables and the omitted individual fixed effect. The partial correlations are the coefficients in the projection of the omitted variable on all included variables.

¹²The expression can be derived in the same way as for Equation 3.3 by redefining the matrices accordingly.

¹³ $M_{[X \ G]}$ is the annihilator matrix for the stacked matrix $[X \ G]$.

force panel include only workers who stay with the same firm; (2) the two surveys refer to different time periods; and (3) the matched data oversample large firms. Brown and Medoff (1989) find support for the view that a substantial part of the size premium can be explained by worker unobservables. However, even after controlling for these, the firm size effect remains significant. This pattern is repeated for Germany (Criscuolo, 2000) and Sweden (Arai, 2003). Similarly, Winter-Ebmer and Zweimüller (1999) and Abowd et al. (1999) find that personal heterogeneity can account for respectively 50% and 70% of the cross-section size premium.

3.1.2 Estimation

I use two main estimation strategies. First, I use a number of panel techniques which treat the explanatory variables as exogenous after controlling for fixed effects. Second, I allow for endogeneity in the explanatory variables, possibly through measurement error or time-varying unobservables, by using Difference- and System-GMM estimators. To simplify the exposition, I stack the x_{it} and s_{it} vectors into a single vector w_{it} and the corresponding coefficient vector ζ .¹⁴ Hence Equation 3.1 is written as

$$y_{it} = w_{it}\zeta + \kappa_i + \varepsilon_{it} \quad (3.5)$$

Equation 3.5 cannot be estimated directly because the number of individual effects increases with the sample size N . First, I estimate a pooled OLS model which is likely to overstate the true size effect as I would expect a positive association between the individual fixed effect and size.¹⁵ Second, I estimate a fixed effects (FE) (or “within”) model, and third a model in first-differences (FD). Both these methods estimate a transformed version of the original equation, eliminating all time-invariant characteristics whether observable or unobservable. The FE estimator uses an equation which is in deviations from the individual-specific mean. Strong exogeneity is a sufficient

¹⁴To be precise, $w_{it} = \{x_{it}s_{it}\}$ is a $[1 \times (P + 1)]$ vector and $\zeta = \begin{Bmatrix} \beta \\ \delta \end{Bmatrix}$ is a $[(P + 1) \times 1]$ vector.

¹⁵Pooled OLS is consistent only if $Cov[u_{it}, w_{it}] = 0$, where $u_{it} = \kappa_i + \varepsilon_{it}$, and either $N \rightarrow \infty$ or $T \rightarrow \infty$.

condition for consistency.¹⁶ The FD estimator transforms the data into differences over time, subtracting the lagged values. The requirement for consistency is somewhat weaker than strong exogeneity.¹⁷ On the other hand, FE is more efficient than FD if ε_{it} is iid. With only two time periods the two estimators are identical. More generally, the fixed effects specification will be preferred because it is more efficient. In all specifications, I account for serial correlation and heteroskedasticity using cluster-robust standard errors.¹⁸

With measurement error in the independent variables or time-varying unobservables, these panel estimators can be inconsistent despite controls for time-invariant fixed effects. In the univariate case, the FD estimator is attenuated as the result of measurement error, i.e. it suffers from an inconsistency toward zero (Arellano, 2003).¹⁹ I use instrumental variable methods to correct for the possible endogeneity. The Difference-GMM (Arellano and Bond, 1991) and System-GMM estimators (Blundell and Bond, 1998) exploit internal instruments in the form of lagged levels or lagged differences of the independent variables to correct for endogeneity in the FD equation.²⁰ With persistent data, System-GMM outperforms Difference-GMM (Blundell and Bond, 1998), but the former requires additional moment conditions which might not hold.²¹

Measurement error in the explanatory variables leads to a contemporaneous correl-

¹⁶The FE estimator is consistent as either $N \rightarrow \infty$ or $T \rightarrow \infty$ and $E[\tilde{\varepsilon}_{it}|\ddot{w}_{it}] = 0$, where $\ddot{w}_{it} = w_{it} - \bar{w}_i$ is the explanatory variable in the transformed model. Strong exogeneity means that $E[\varepsilon_{it}|w_{i1}, \dots, w_{iT}] = 0$.

¹⁷Consistency requires that $E[\Delta\varepsilon_{it}|\Delta w_{it}] = 0$, where $\Delta w_{it} = w_{it} - w_{i,t-1}$.

¹⁸I cluster at the individual level, which is appropriate for panel data (Arellano, 1987).

¹⁹Compared with the equation in levels, the FD estimator might exacerbate the inconsistency arising from measurement error, if the true (unobserved) explanatory variable is a stationary time series with positive serial correlation. In the multivariate case with only one explanatory variable measured with error, the coefficient on that variable is attenuated, and the coefficients on all other variables are inconsistent. With measurement error in more than one variable, the direction of the inconsistency becomes very hard to establish.

²⁰Note that I cannot use the fixed effects estimator, because the within transformation introduces all past values.

²¹Instruments are valid if they are uncorrelated with the error term, as discussed in more detail below. Furthermore, instruments need to be informative, i.e. they need to be correlated with the endogenous variable. If the internal instruments are informative for enterprise size, it implies that size does not follow Gibrat's law. This law states that the growth rate, i.e. the first-difference of a variable in logs, and the level are independent. The empirical evidence overwhelmingly rejects Gibrat's law (Sutton, 1997).

ation between w_{it} and ε_{it} , and between Δw_{it} and $\Delta \varepsilon_{it}$ in the FD equation. Under weak exogeneity, I can use lagged levels as instruments for contemporaneous differences, as exploited by the Difference-GMM estimator.²² In addition, the System-GMM estimator uses lagged differences as instruments for contemporaneous levels, which requires the further assumption that changes in the explanatory variable are uncorrelated with the individual fixed effect κ_i .²³ I use a two-step GMM estimator with standard errors which are corrected for small sample bias using the (Windmeijer, 2005) procedure and which are robust to serial correlation and heteroskedasticity.²⁴ I test for the robustness of the results to using different instrument sets.

3.2 Data description

I use two datasets to estimate the effect of enterprise size on wages. The first is a matched firm-worker dataset (covering the period 1995-2002) which provides a consistent estimate of the size effect in wage-employment since I can control for unobserved person and firm effects. The second is a labour force dataset (covering the period 2004-2008) which provides comparable data for the self-employed and serves as a robustness check for the wage-employed.

The first data (Ghana RPED/GMES) are a panel survey of manufacturing firms in Ghana and, in addition, a sample of their workers (Teal, 2000).²⁵ The data cover the four major urban areas of Ghana – Greater Accra, Kumasi, Takoradi and Cape Coast.

²²If I assume that past values of the explanatory variable are uncorrelated with the error term, so w_{it} is weakly exogenous apart from being contemporaneously correlated with ε_{it} , then $E[w_{is}\varepsilon_{it}] = 0$ for $s \leq t - 1$. This implies that in the first-differences equation $E[w_{is}\Delta\varepsilon_{it}] = 0$ for $s \leq t - 2$ holds. I also need to assume serially uncorrelated errors and the following error components assumption: $E[\kappa_i] = E[\varepsilon_{it}] = E[\kappa_i\varepsilon_{it}] = 0$.

²³Under the additional assumption that $E[\Delta w_{is}\kappa_i] = 0$ for $s = 2, \dots, T$, weak exogeneity implies that in the levels equation the condition $E[\Delta w_{is}u_{it}] = 0$ for $s \leq t - 1$ holds, where $u_{it} = \kappa_i + \varepsilon_{it}$, as before.

²⁴Asymptotically, the two-step estimator is more efficient than the one-step estimator. However, the asymptotic two-step standard errors are downward biased in small samples. As Bond and Windmeijer (2002) show, the standard errors proposed by Windmeijer (2005) lead to more reliable results than the asymptotic standard errors.

²⁵The first three waves were administered by the World Bank's Regional Program on Enterprise Development (RPED). Rounds four to seven were collected by the Centre for the Study of African Economies (CSAE) at the University of Oxford, the University of Ghana at Legon and the Ghana Statistical Office. These rounds are referred to as the Ghana Manufacturing Enterprise Survey (GMES).

The original firm sample incorporated 200 firms, based on the 1987 Ghana Census of Manufacturing Activities.²⁶ I use the last four waves of the data collected between 1995 and 2003, which provide data spanning the period 1995 to 2002.²⁷ Each of these surveys covers a two-year period with the exception of the last survey which covers three years. Time-varying characteristics in the intervening years are collected using recall questions.²⁸ I exploit the panel structure introduced by the recall questions, and match workers across waves where possible.²⁹ For each firm, a sample of up to ten workers were interviewed, that were chosen to be representative of the firm’s workforce in terms of its occupational structure.³⁰ This sampling procedure leads to workers in large firms being under-sampled relative to workers in small firms. However, the oversampling of large firms counteracts this effect. In sum, it is not straightforward to extrapolate these results to other contexts. The results from the labour force panel serve as a robustness check.

The second dataset (the Ghana Urban Panel Survey) is a labour force survey covering a representative sample of the urban labour force in the same urban areas as the firm survey (Teal, 2008; Falco et al., 2011).³¹ I use the first four waves of the survey collected between 2004 and 2008, which cover five years of data.³² The survey collects comparable data across all employment categories, including the self-employed. Meas-

²⁶Large enterprises were oversampled in the original survey. It has not been possible to obtain the necessary census information to construct sampling weights. Sample attrition has been relatively moderate, especially for a developing country dataset. Dropouts were replaced by similar firms in terms of size and sector.

²⁷Individual wages coincide with the contemporaneous period (i.e. 1995 to 2003) while the firm questions refer to the previous year (i.e. 1994 to 2002). Hence I lose one year when I match the two parts of the data. The first three waves were annual surveys, which I exclude because they do not have the same structure using recall questions.

²⁸Recall bias could be an issue. However, using recall questions helps to reduce the problem arising from attrition which may be more severe in a developing country dataset.

²⁹For most individuals I only have two or three time-series observations. This is because the matched firm-worker data, in contrast to the labour force panel, were not designed as a panel survey of workers from the outset, although efforts were made to re-interview the same worker.

³⁰Apprentices were interviewed separately and I exclude them from the analysis, since they typically receive very low wages in Ghana.

³¹The data collection was directed by the CSAE and supported by the Ghana Statistical Office. The initial sample was a stratified random sample from the 2000 census.

³²The 2008 wave includes recall questions for 2007. In 2004, some of the firm workers from the GMES (the matched firm-worker survey) were re-surveyed as part of the labour force survey. I focus on the representative sample and exclude these additional observations, following Rankin et al. (2010). However, the results are robust to including these observations.

uring the earnings of the self-employed, especially in a developing country context, is difficult.³³ Rankin et al. (2010) review in detail the survey instruments that were used to make these data as reliable as possible (e.g. the use of electronic data collection). The earnings data for the self-employed have been used by a number of papers which included some robustness checks (Sandefur et al., 2007; Rankin et al., 2010; Falco et al., 2011; Falco, 2014).

In both datasets, the sample used in the estimation is restricted to working-age individuals, defined as being between 15 and 65 years of age.³⁴ Where necessary, I average supposedly time-invariant variables (e.g. gender, education) over time and make time-deterministic (e.g. age and tenure) variables consistent.³⁵ In the matched firm-worker survey, I use monthly earnings after tax including all benefits. In the labour force survey, I use before-tax earnings (monthly, including benefits) which is the only earnings variable I have access to in the data.³⁶ In wage-employment, firm size refers to the number of employees in the firm which would typically exclude the entrepreneur.³⁷ In self-employment, the number of employees includes the entrepreneur, so a self-

³³For example, the self-employed often lack written accounts or the distinction between business and personal expenditures can be difficult for household-based businesses.

³⁴In the matched firm-worker data, I exclude the top and bottom 1% of real earnings to account for potential measurement error. In addition, I excluded three firms listed in two locations as I include location controls in the individual-level regressions.

³⁵Education, or schooling, is treated as a time-invariant characteristic once the individual has entered the labour force, i.e. is included in the data. As a result of measurement error, there could exist time variation in variables which ought to be time-invariant. I average these time-invariant variables over time and round dummy variables where necessary. The latter was only necessary for very few observations. Furthermore, I ensure that the difference in the age and tenure variables between two consecutive years is exactly one.

³⁶Furthermore, the labour force survey collects comparable earnings data for the wage- and self-employed. Before-tax earnings might be more comparable because it is difficult to measure the taxes paid by the self-employed in Ghana. Falco et al. (2011) also use before-tax earnings. Using before-tax earnings might bias the size coefficient if workers in larger firms are more likely to pay taxes. In the matched firm-worker survey, I observe earnings both before and after taxes. I find very similar coefficients on firm size when using earnings before tax compared to the after tax estimates in the main results. The coefficient estimates are 0.181 compared with 0.166 (Table 3.3, column 1), and 0.175 compared with 0.171 (Table 3.3, column 4) when controlling for physical capital. The coefficient lies within 1.28 standard errors (0.25 standard errors when controlling for capital) of the original coefficient in Table 3.3.

³⁷I use the number of employees as the main size measure because (1) it is measured quite reliably, which is particularly important in a developing country context, and (2) it is most commonly used in the literature (e.g. Oi and Idson, 1999). As a robustness check, I included the physical capital stock in the matched firm-worker data where this is available. Firm size refers to the size of the establishment. In US data, Brown and Medoff (1989) find that the effect of establishment size on wages is larger than that of company size. The firms in my data are unlikely to have multiple establishments. None of the self-employed individuals have multiple establishments.

employed individual with no employees works in an establishment of size one.

3.2.1 Summary statistics

Table 3.1 presents summary statistics for the matched firm-worker survey split according to the firm size categories used by Söderbom and Teal (2004). Employees in large firms earn approximately twice as much per month as employees in the smallest size category, where labour contracts are usually absent.³⁸ However, there exists substantial heterogeneity around these average wages (Figure 3.1). On average, workers in larger firms are older, more educated, have more experience, and are more likely to be men (in particular when moving from micro to small firms). There is no clear trend in average hours, but they are perhaps measured with some error and hard to interpret as most employees have little choice over their hours.³⁹

Table 3.2 shows the summary statistics for the labour force survey. The self-employed tend to be older and less educated, and are more likely to be women (especially amongst those without employees). On average, earnings in wage- and self-employment are quite similar, particularly amongst the self-employed with employees.⁴⁰ However, the earnings of the self-employed are more spread out (Figure 3.2), and in terms of the median there seems to be a premium for the wage-employed over the self-employed. Nevertheless, the substantial overlap between the sectors suggests that self-employment is not always an inferior outcome, at least in terms of the level of earnings.⁴¹ Size increases earnings within both wage- and self-employment although

³⁸In 2007, Ghana re-valued its currency by 10,000:1. I use new Cedis throughout.

³⁹Therefore, I follow Falco et al. (2011) and control for hours in the earnings regressions instead of using hourly earnings as the dependent variable.

⁴⁰Earnings of the self-employed are self-reported profits net of routine operating expenses and gross of fixed capital expenditure (Rankin et al., 2010).

⁴¹However, self-employment earnings are probably more volatile. In the original Harris and Todaro (1970) model, an informal self-employment sector arises as the result of a formal sector which pays above market-clearing wages (and in the absence of open unemployment). Traditionally, urban self-employment has been viewed as the undesirable, or “murky” (Fields, 1975), reserve option for those who queue up for a formal sector job. The more recent literature recognises the heterogeneity within the informal sector and the substantial overlap in earnings (Maloney, 2004; Sandefur et al., 2007; Teal, 2011). Fajnzylber et al. (2006) suggest that the informal sector presents desirable job opportunities to at least some workers. The share of earnings inequality (e.g. according to GE(0)) explained by differences between the wage- and self-employed is minimal compared to the inequality that exists within these occupations.

there is again a substantial overlap (Figure 3.2). The earnings of employees in small firms and sole traders are quite similar.

As already mentioned, the matched data oversample larger firms compared with the representative labour force survey, as shown by comparing the respective firm size distributions (Figure 3.A.3). Hence the two datasets are informative about workers in different parts of the firm size distribution. There exists substantial cross-sectional dispersion in firm size in both surveys, as shown by Table 3.A.1 (for the matched data) and Table 3.A.3 (for the labour force survey). This is important for the identification of the size coefficient in the pooled cross-section, while the panel regressions rely on variation over time. Firm size appears to be more persistent in the labour force data (Table 3.A.4) than the matched data (Table 3.A.2), particularly in 2008. In self-employment (Tables 3.A.5 and 3.A.6), both the cross-sectional and the temporal variation are more limited than for wage-employees. Of course, part of these effects could also be driven by attrition, which is a concern in any panel dataset especially from developing countries. Using the same labour force data, Falco et al. (2011) correct for attrition and find that it does not bias their results.

3.3 Results

3.3.1 The effect of scale in wage-employment

In this section I estimate the size effect for the wage-employed using two datasets. First, I use the matched employer-employee data which allow me to control for both person- and firm-unobservables and thus obtain a consistent estimate of the size effect (see Section 3.1.1). However, these data are not representative and they only include workers who stay with the same firm. Therefore, I test for the robustness of the results by using a labour force survey in the second part. While these data are representative of the urban labour force, the estimated size coefficient might be inconsistent due to omitting firm unobservables.

3.3.1.1 Results from the matched data

I briefly mention the other variables in Table 3.3 before discussing the size coefficient in more detail. In the pooled cross-section, an individual working 42 hours per week earns approximately 1% more than an individual working 40 hours. All else equal, wages are 16% higher for men than for women. The cross-section returns to education are convex, which is consistent with the recent literature on developing countries (Söderbom et al., 2006).⁴² The age, tenure and year dummies are difficult to interpret in the FE and FD equations, because the transformations introduce perfect multicollinearity.

In the pooled cross-section I find a highly significant positive relationship between firm size and wages. An individual in a firm which is 1% larger earns 0.17% more per month which compares with 0.21% in Söderbom et al. (2005). Using the FE estimator, the coefficient estimate drops by two thirds and it remains significant although at a lower level. In the FD specification, the size effect is even smaller and insignificant, which is not driven by differences in the standard errors.⁴³ I check for the robustness of these results to separate analyses for men and women. The size effect in the cross-section is somewhat larger for women (Table 3.A.8). However, after controlling for fixed effects, there remains only a highly significant effect for men (Table 3.A.7).

In a further robustness check I include a control for physical capital which might capture other aspects of scale.⁴⁴ Because firm size and physical capital are highly collinear, their coefficients are difficult to interpret and the true size effect is a combination of the two coefficient estimates.⁴⁵ Including the capital control leaves the results for the pooled OLS and the FE estimator almost unchanged.⁴⁶ However, in the FD model, physical capital is now significant which is driven by the women in the sample. In the

⁴²The theoretical model of schooling choice by Card (1999) suggests a concave education-earnings profile. Recent evidence suggests convexity in at least some parts of the domain of the earnings function (Belzil and Hansen, 2002).

⁴³The smaller estimate in FD compared with FE might be due to measurement error, which leads to a larger bias in the FD than the FE estimator.

⁴⁴For example, changes in the capital stock might capture unobserved changes in the skill composition of the workforce as a result of capital-skill complementarities.

⁴⁵The correlation coefficient between firm size and physical capital is 0.841.

⁴⁶There is no evidence that the physical capital effect is due to the capital-labour ratio. If the effect were due to the capital-labour ratio, the coefficients on $\ln(\text{physical capital})$ and $\ln(\text{firm size})$ would have to be equal in absolute value and of opposite sign, with the coefficient on capital being positive.

FE model, firm size is only statistically significant for men, and this is robust to the capital control. In sum, there is strong evidence for a size effect in the FE and some mixed results in the FD models. In the FE estimation, the size effect works through the number of employees, and is driven by the male sub-sample. In the FD model, it operates through the physical capital variable and is driven by the women in the sample.

As argued above, these estimated coefficients might be inconsistent for example as a result of measurement error or time-varying unobservables. I use Difference- and System-GMM (Table 3.4) estimators to correct for this, reporting only the coefficients on firm size and the capital stock.⁴⁷ I find larger coefficient estimates, but the standard errors also increase, as is typically the case when using instrumental variables. I find no significant effect of size on wages using a Difference-GMM estimator which is robust to restricting the instrument set to the $t - 3$ dated instruments and the inclusion of the capital control.⁴⁸ The System-GMM estimates are of a similar order of magnitude as the pooled cross-section. They are significant when using the restricted set of instruments and with physical capital, but this is not robust to using the full (column 3) and collapsed (columns 5 and 10) instrument sets. The method of collapsing the instrument set is suggested by Roodman (2009a).

3.3.1.2 Results from the labour force data

I replicate the analysis using data on wage-employees from the labour force data (Table 3.5). As before, monthly earnings increase with the number of hours worked and the returns to education are convex. I find a strongly significant size effect in the

⁴⁷I report the p-values of the Hansen test for over-identifying restrictions which suggest that the set of instruments is valid. Furthermore, the p-values are not too close to 1 which would point towards a Hansen test weakened by instrument proliferation (Roodman, 2009b). I also test the null hypothesis of no serial correlation in the first-differences residuals (Arellano and Bond, 1991). The test for first-order serial correlation ($AR(1)$) rejects the null because $\Delta\varepsilon_{it}$ and $\Delta\varepsilon_{i,t-1}$ have a common term in $\varepsilon_{i,t-1}$. A test for first-order serial correlation in the levels equation requires a test for second-order serial correlation in the first-differences residuals. Because all p-values of the $AR(2)$ test exceed 0.05, I fail to reject the null hypothesis of no serial correlation at the 5% level.

⁴⁸I also collapsed the instrument set using the method suggested by Roodman (2009a) (results are not reported in the table). In this specification, the coefficient estimate is also insignificant and the Hansen test suggests that the instrument set is not valid (p-value=0.055). I chose to use the $t - 3$ instruments, because the instrument set is invalid when using $t - 2$ instruments.

pooled cross-section which is smaller than in the matched data.⁴⁹ I find a statistically significant effect of firm size on wages in both the FE and the FD equations. Firm size is only significant for men, but not for women (Table 3.A.9). These results are very similar to the estimates from the matched dataset.⁵⁰ Table 3.6 shows the results of the Difference- and System-GMM in an attempt to account for endogeneity possibly arising from measurement error or omitting the firm fixed effects in the labour force data. There is no longer any evidence of a significant effect of firm size on wages, which is robust to changes in the instrument set and using instruments which pass the standard diagnostics.⁵¹

3.3.2 The size effect in self-employment

Table 3.7 estimates the effect of enterprise size on earnings in self-employment using the labour force data. Many of the covariates have similar effects compared with the wage-employed, except that education does not appear to have a significant effect on the earnings of the self-employed. I find a large cross-section size premium in self-employment, which is even larger than the cross-section premium for wage-employees above.⁵² This size premium appears to be driven by time-invariant unobservables, as it is no longer significant after controlling for these characteristics in the FE and FD specifications. This finding is robust to estimating separately by gender (Table 3.A.10). This contrasts with the results for wage-employees where firm size remains significant even after including these controls.

The Hansen test for the Difference-GMM specification (Table 3.8) suggests that the instrument set is invalid. This could be due to the persistence in enterprise size

⁴⁹The coefficient estimates are somewhat smaller than in Sandefur et al. (2007) and Falco et al. (2011) who use the same data. However, these authors use different years and estimate a pooled regression including both wage- and self-employed. Falco et al. (2011) confirm the finding of a significant size effect in the cross-section and FE specifications but not the System-GMM.

⁵⁰However, the results from the two datasets are hard to compare for the reasons discussed earlier.

⁵¹The tests for serial correlation in the error terms are all fine. The p-values of the Hansen test all suggest that the set of instruments is valid. However, the value observed in the Difference-GMM with the full set of instruments is approaching 1, which might be a sign of instrument proliferation (Roodman, 2009b).

⁵²The coefficient on the number of employees has a similar magnitude to the estimates in Falco et al. (2011) and Sandefur et al. (2007), although these authors pool the wage- and self-employed in the same regression.

over time (Table 3.A.6), as the Difference-GMM estimator exploits this variation.⁵³ The lack of a significant size effect in the FE and FD estimators could perhaps be explained by the same issue. The set of instruments used in the System-GMM, which additionally uses the levels equation, is valid. There is no evidence of a significant size effect in this estimation. In sum, these results suggest that the strong cross-sectional relationship between enterprise size and earnings in self-employment is driven by time-invariant unobservables or endogeneity.⁵⁴ From a policy perspective, it matters whether these unobservables are inborn characteristics or whether they can be learnt through education or training.

3.4 Conclusion

This chapter has attempted to understand the sources of the cross-sectional enterprise size premium observed in both self- and wage-employment in Ghana. In both sectors, a substantial size premium remains after controlling for observable characteristics, including human capital. For the self-employed, this premium seems to be driven by time-invariant unobservables, which could be interpreted as entrepreneurial ability. On the other hand, for the wage-earners, there is some evidence of a size premium even after controlling for time-invariant unobservables.⁵⁵

While I used the best available panel data and tested for robustness using a number of estimation methods, the estimates could of course be invalid for a number of reasons. Measuring earnings, particularly those of the self-employed, is difficult in developing countries. The panel element of the data allows me to control for time-invariant characteristics in the error term whose presence could make the size variable endogenous. However, these panel data estimators rely on variation over time for identification. This variation may be limited over the observation period, particularly for

⁵³Blundell and Bond (1998) show that the System-GMM performs better than the Difference-GMM with persistent data. This is of course under the assumption that the additional moment conditions of the System-GMM are met.

⁵⁴This is consistent with the Lucas (1978) model of the firm size distribution, in which individuals have a fixed, unobservable entrepreneurial talent.

⁵⁵For employees, I cannot disentangle the time-invariant unobservable into person- and firm-effects. Hence the effect could be driven by a combination of the two.

self-employed enterprises.

Understanding the determinants of earnings has important implications for poverty reduction and economic development in Africa, particularly in urban areas.⁵⁶ The most important asset of the poor is their labour, so improving the returns to labour is paramount (Teal, 2011). The World Development Report 2013 on Jobs is one indication that the development policy community has recognised this issue (World Bank, 2012). The report stresses the quality of jobs, one dimension of which is the wage paid. In many African countries, including Ghana, outright unemployment is low, as the informal sector absorbs this labour. In fact, urban self-employment is the fastest growing type of employment in Ghana. Some authors view this as a sign of a vibrant labour market, providing desirable job opportunities to some individuals (e.g. Fajnzylber et al., 2006). Other authors argue that this is the result of an underperforming economy where firms fail to scale up and provide sufficient employment (e.g. Fields (1975), see discussion on p. 113).

As this chapter has shown, earnings in self-employment and small firms are often low. Wages increase with the size of the firm. Of course, these estimates cannot be interpreted as causal. Instead, larger firms possess some other characteristic that is associated with higher wages.⁵⁷ This chapter has attempted to provide a starting point to an analysis of these mechanisms.

⁵⁶In rural areas, agriculture is still the dominant economic activity. It is important to bear in mind that poverty measurement in Africa typically uses consumption expenditure (see the Chapter on global inequality), which is distinct from labour market earnings.

⁵⁷According to Melitz (2003) more productive forms are more likely to export, as there are fixed costs of exporting, and they are also bigger. Furthermore, the size of firms that do not export is constrained by the size of the domestic market. Söderbom and Teal (2004) find that larger firms in Ghana have lower costs of capital, so they may be more likely to use more productive technologies. In cross-country evidence, Bloom et al. (2010) find that financial constraints appear to limit firm growth mostly in small firms, but less so in larger firms. They also suggest that firm growth is constrained by limited delegation in firms in developing countries, perhaps as a result of a lack of trust.

3.5 Main Tables and Figures

Table 3.1: Matched firm-worker data: Summary statistics by firm size

	Micro Size < 6	Small Size 6-29	Medium Size 30-99	Large Size > 99	Total
Real monthly earnings, 1997 prices	12 (8)	12 (9)	17 (13)	25 (19)	18 (13)
Nominal monthly earnings, USD	54 (36)	48 (39)	75 (56)	111 (82)	79 (55)
Hours per week	53 (52)	46 (44)	43 (40)	45 (40)	45 (40)
Share of men (percent)	59 (100)	74 (100)	84 (100)	88 (100)	82 (100)
Age	32 (30)	33 (30)	37 (35)	40 (39)	37 (35)
Tenure	7 (4)	8 (6)	8 (6)	9 (7)	8 (6)
Education	9 (10)	10 (10)	12 (10)	13 (12)	11 (10)
Firm size	4 (4)	17 (16)	57 (54)	307 (206)	126 (56)
Capital-labour ratio	85 (21)	140 (51)	339 (166)	692 (317)	391 (176)
Observations	100	1715	2122	1925	5862

Main statistic: mean. In parentheses: median. Capital-labour ratio is in 1991 prices.

Figure 3.1: Matched firm-worker data: Distribution of earnings by size




Table 3.2: Labour force data: Summary statistics by occupation and size

	Wage: Size<30	Wage: Size>30	Self: Sole-Trader	Self: Employees	Total
Real monthly earnings, 1997 prices	17 (12)	40 (26)	23 (11)	38 (18)	26 (14)
Nominal monthly earnings, USD	76 (56)	190 (128)	108 (48)	189 (84)	122 (64)
Hours per week	48 (48)	46 (41)	48 (50)	51 (50)	48 (48)
Share of men (percent)	60 (100)	73 (100)	28 (0)	43 (0)	44 (0)
Age	31 (28)	36 (34)	36 (35)	39 (39)	35 (34)
Tenure	6 (4)	10 (6)	10 (7)	12 (10)	9 (7)
Education	9 (9)	11 (11)	7 (9)	8 (9)	8 (9)
Firm size	7 (7)	129 (75)	. (.)	. (.)	55 (15)
No. Employees	. (.)	. (.)	1 (1)	4 (3)	1 (1)
Observations	856	556	1732	408	3552

Main statistic: mean. In parentheses: median.

Figure 3.2: Labour force data: Distribution of earnings by size and sector




Table 3.3: Matched firm-worker data: Baseline results

Ln(Real earnings)	(1) POLS	(2) FE	(3) FD	(4) POLS	(5) FE	(6) FD
Ln(Firm size)	0.166*** (0.0117)	0.0500* (0.0219)	0.0207 (0.0195)	0.171*** (0.0164)	0.0495* (0.0219)	0.0208 (0.0196)
Ln(Capital)				-0.00293 (0.00841)	-0.0135 (0.0270)	0.0380** (0.0140)
Ln(Hours)	0.217*** (0.0461)	0.135* (0.0676)	0.0793 (0.0617)	0.215*** (0.0461)	0.135* (0.0677)	0.0792 (0.0617)
Male	0.155*** (0.0299)			0.156*** (0.0300)		
Age	0.0245*** (0.00681)	0.122*** (0.0218)		0.0242*** (0.00686)	0.122*** (0.0220)	
Age squared/100	-0.0147 (0.00849)	-0.0860** (0.0321)	-0.102*** (0.0295)	-0.0144 (0.00854)	-0.0853** (0.0320)	-0.0972** (0.0298)
Tenure	0.00643** (0.00247)			0.00646** (0.00249)		
Tenure squared/100	-0.00781 (0.00398)	0.0134 (0.0383)	-0.00210 (0.0397)	-0.00792* (0.00398)	0.0130 (0.0378)	-0.00308 (0.0399)
Education	0.000956 (0.00735)			0.00122 (0.00739)		
Education squared/100	0.268*** (0.0358)			0.267*** (0.0360)		
Observations	5862	5862	3530	5835	5835	3500

Standard errors in parentheses

Dependent variable: Ln(Real monthly earnings, 1997 prices). Standard errors clustered on the individual.

Including year dummies, and controls for location, sector, public enterprise and foreign ownership.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.4: Matched firm-worker data: Difference- and System-GMM

Ln(Real earn.)	(1) D1	(2) D2	(3) S1	(4) S2	(5) S3	(6) D1	(7) D2	(8) S1	(9) S2	(10) S3
Ln(Firm size)	0.0856 (0.135)	0.0924 (0.0844)	0.0943 (0.0732)	0.186** (0.0703)	-0.126 (0.181)	0.0836 (0.127)	0.0842 (0.0811)	0.133* (0.0615)	0.274** (0.0970)	0.0127 (0.159)
Ln(Capital)						-0.0754 (0.0515)	-0.0122 (0.0759)	-0.0166 (0.0371)	-0.00579 (0.0568)	-0.0123 (0.0692)
Observations	3504	3504	5862	5862	5862	3474	3474	5835	5835	5835
Individuals	1842	1842	2325	2325	2325	1828	1828	2325	2325	2325
Instruments	193	52	235	89	81	214	57	262	99	88
Hansen	0.156	0.379	0.312	0.348	0.315	0.234	0.357	0.347	0.205	0.279
AR1	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
AR2	0.119	0.165	0.121	0.146	0.127	0.177	0.201	0.152	0.170	0.154

Standard errors in parentheses

Specification as in Table 3.3. Columns 1,2,6,7: Difference-GMM, remainder is System-GMM. Two-step GMM estimator.

Columns 1,3,6,8: Full set of instruments. Columns 2,4,7,9: Only lags t-3. Columns 5,10: Collapsed.

All variables instrumented except year, location, sector and ownership controls.

p-values of a (robust) Hansen test of overidentifying restrictions and for AR(1) and AR(2) residuals reported.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.5: Labour force data, wage-employees: Baseline results

Ln(Real earnings)	(1) POLS	(2) FE	(3) FD
Ln(Firm size)	0.119*** (0.0138)	0.0531** (0.0186)	0.0350* (0.0155)
Ln(Hours)	0.181* (0.0785)	0.249** (0.0901)	0.274** (0.0937)
Male	0.162** (0.0507)		
Age	0.0965*** (0.0172)		
Age squared/100	-0.109*** (0.0227)	-0.255* (0.121)	-0.201 (0.121)
Tenure	-0.00233 (0.0118)	0.309*** (0.0727)	
Tenure squared/100	0.0541 (0.0426)	-0.0513 (0.161)	0.0943 (0.138)
Education	-0.0284 (0.0272)		
Education squared/100	0.594*** (0.149)		
Observations	1412	1412	600

Standard errors in parentheses

Dependent variable: Ln(Real monthly earnings, 1997 prices. Standard errors clustered on the individual. Including year dummies, and controls for location and public ownership.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.6: Labour force data, wage-employees: Difference- and System-GMM

Ln(Real earnings)	(1) D1	(2) D2	(3) D3	(4) S1	(5) S2	(6) S3
Ln(Firm size)	0.0939* (0.0448)	0.0890 (0.0597)	0.264 (0.397)	0.0822* (0.0414)	0.0871 (0.0587)	-0.0424 (0.129)
Observations	600	600	600	1412	1412	1412
Individuals	355	355	355	729	729	729
Instruments	59	32	23	78	41	37
Hansen	0.758	0.700	0.742	0.506	0.521	0.665
AR1	0.011	0.024	0.039	0.011	0.014	0.009
AR2	0.108	0.096	0.162	0.125	0.131	0.109

Standard errors in parentheses

Specification as in Table 3.5. Difference-GMM (columns 1-3) and System-GMM (columns 4-6)

Columns 1,4: All instruments. Columns 2,5: Collapsed. Columns 3,6: only t-3. Two-step GMM estimator.

All variables instrumented except year, location and ownership controls.

p-values of a (robust) Hansen test of overidentifying restrictions and for AR(1) and AR(2) residuals reported.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.7: Labour force data, self-employed: Baseline results

Ln(Real earnings)	(1) POLS	(2) FE	(3) FD
Ln(No. Employees)	0.264*** (0.0608)	0.0629 (0.0847)	0.0739 (0.0844)
Ln(Hours)	0.299*** (0.0609)	0.240*** (0.0635)	0.0873 (0.0719)
Male	0.376*** (0.0622)		
Age	0.0727** (0.0226)		
Age squared/100	-0.0833** (0.0289)	0.0917 (0.177)	-0.0683 (0.176)
Tenure	0.0257* (0.01000)	0.129 (0.101)	
Tenure squared/100	-0.0342 (0.0277)	-0.0350 (0.223)	0.166 (0.248)
Education	0.0400 (0.0240)		
Education squared/100	-0.158 (0.182)		
Observations	2140	2140	1061

Standard errors in parentheses

Dependent variable: Ln(Real monthly earnings, 1997 prices. Standard errors clustered on the individual. Including year dummies, and controls for location.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.8: Labour force data, self-employed: Difference- and System-GMM

Ln(Real earnings)	(1) D1	(2) D2	(3) D3	(4) S1	(5) S2	(6) S3
Ln(No. Employees)	0.635 (0.402)	2.068** (0.674)	2.378** (0.751)	0.0330 (0.249)	0.416 (0.449)	0.00267 (0.918)
Observations	1061	1061	1061	2140	2140	2140
Individuals	583	583	583	950	950	950
Instruments	58	31	22	77	40	36
Hansen	0.006	0.016	0.026	0.674	0.632	0.510
AR1	0.000	0.000	0.000	0.000	0.000	0.000
AR2	0.315	0.374	0.381	0.390	0.283	0.300

Standard errors in parentheses

Specification as in Table 3.7. Difference-GMM (columns 1-3) and System-GMM (columns 4-6)

Columns 1,4: All instruments. Columns 2,5: Collapsed. Columns 3,6: only t-3. Two-step GMM estimator.

All variables instrumented except year and location controls.

p-values of a (robust) Hansen test of overidentifying restrictions and for AR(1) and AR(2) residuals reported.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

3.A Appendix Tables and Figures

Figure 3.A.1: GDP and firm size across countries




Figure 3.A.2: GDP and wage employment across countries




Figure 3.A.3: Comparing distribution of firm size in matched and labour force data




Table 3.A.1: Matched firm-worker data: Distribution of firm size by year

Firm size	1995	1996	1997	1998	1999	2000	2001	2002	Total
Size≤2	0.12	0.37	0.00	0.63	0.00	0.30	0.00	0.00	0.20
2<Size≤10	8.21	8.67	7.71	9.38	10.11	8.73	8.59	9.95	8.85
10<Size≤20	13.04	15.14	12.02	9.27	9.03	14.05	14.32	14.43	12.28
20<Size≤50	25.85	23.81	19.84	25.40	25.03	29.14	29.95	26.87	25.09
50<Size≤100	17.39	19.41	25.28	20.44	20.58	19.23	20.05	25.37	20.86
100<Size	35.39	32.60	35.15	34.88	35.26	28.55	27.08	23.38	32.70
Total	100.00 (828)	100.00 (819)	100.00 (973)	100.00 (949)	100.00 (831)	100.00 (676)	100.00 (384)	100.00 (402)	100.00 (5862)

Number of observations in parentheses.

Table 3.A.2: Matched firm-worker data: Change in firm size since last year

Changes in firm size	1996	1997	1998	1999	2000	2001	2002
No change	8.79 (64)	38.41 (63)	5.79 (50)	25.06 (103)	3.79 (23)	31.70 (110)	27.42 (105)
Increase<10%	22.12	12.80	12.96	14.84	13.67	24.21	11.49
Increase≥10%	20.33	18.29	32.87	15.82	32.78	8.93	17.49
Decrease<10%	10.58	15.24	16.78	19.95	15.65	17.87	18.54
Decrease≥10%	38.19	15.24	31.60	24.33	34.10	17.29	25.07
Total	100.00 (728)	100.00 (164)	100.00 (864)	100.00 (411)	100.00 (607)	100.00 (347)	100.00 (383)

Number of observations in parentheses.

Table 3.A.3: Labour force data: Distribution of firm size by year

Firm size	2004	2005	2006	2007	2008	Total
Size $\leq$ 2	19.91	9.25	24.23	22.97	23.89	20.33
2<Size $\leq$ 10	36.20	43.06	20.70	11.96	17.00	25.64
10<Size $\leq$ 20	7.24	10.68	20.93	16.27	12.96	14.66
20<Size $\leq$ 50	8.14	10.32	14.32	18.66	20.24	14.24
50<Size $\leq$ 100	8.60	7.12	4.19	10.05	8.91	7.15
100<Size	19.91	19.57	15.64	20.10	17.00	17.99
Total	100.00 (221)	100.00 (281)	100.00 (454)	100.00 (209)	100.00 (247)	100.00 (1412)

Number of observations in parentheses.

Table 3.A.4: Labour force data: Change in firm size since last year

Changes in firm size	2005	2006	2007	2008
No change	49.17 (59)	0.00 (0)	40.19 (43)	96.70 (176)
Increase $<10\%$	0.00	9.95	0.00	0.00
Increase $\geq 10\%$	28.33	40.84	35.51	2.75
Decrease $\geq 10\%$	22.50	49.21	24.30	0.55
Total	100.00 (120)	100.00 (191)	100.00 (107)	100.00 (182)

Number of observations in parentheses.

Table 3.A.5: Labour force data: Distribution of number of employees by year

No. Employees	2004	2005	2006	2007	2008	Total
Size=1	90.30	85.11	85.37	67.10	68.53	80.93
Size=2	2.99	4.23	7.31	18.24	17.94	8.97
Size=3	3.48	3.31	5.12	7.49	6.47	4.91
Size ≥ 4	3.23	7.35	2.19	7.17	7.06	5.19
Total	100.00 (402)	100.00 (544)	100.00 (547)	100.00 (307)	100.00 (340)	100.00 (2140)

Number of observations in parentheses.

Table 3.A.6: Labour force data: Change in no. employees since last year

Changes in No. Employees	2005	2006	2007	2008
No change	79.39 (208)	77.17 (267)	62.65 (104)	98.61 (283)
Increase $\geq 10\%$	13.74	8.96	30.12	0.70
Decrease $\geq 10\%$	6.87	13.87	7.23	0.70
Total	100.00 (262)	100.00 (346)	100.00 (166)	100.00 (287)

Number of observations in parentheses.

Table 3.A.7: Matched firm-worker data: Analysis for men only

Ln(Real earnings)	(1) POLS	(2) FE	(3) FD	(4) POLS	(5) FE	(6) FD
Ln(Firm size)	0.161*** (0.0127)	0.0714** (0.0224)	0.0365 (0.0203)	0.194*** (0.0187)	0.0704** (0.0224)	0.0355 (0.0205)
Ln(Capital)				-0.0224* (0.00972)	-0.0237 (0.0294)	0.0238 (0.0146)
Observations	4805	4805	2870	4780	4780	2843

Standard errors in parentheses

Specification as in Table 3.3, but only male workers are included in this table.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.A.8: Matched firm-worker data: Analysis for women only

Ln(Real earnings)	(1) POLS	(2) FE	(3) FD	(4) POLS	(5) FE	(6) FD
Ln(Firm size)	0.180*** (0.0282)	-0.0210 (0.0544)	-0.0541 (0.0481)	0.0930** (0.0349)	-0.0184 (0.0535)	-0.0482 (0.0481)
Ln(Capital)				0.0630*** (0.0178)	0.00897 (0.0580)	0.0699* (0.0341)
Observations	1057	1057	660	1055	1055	657

Standard errors in parentheses

Specification as in Table 3.3, but only female workers are included in this table.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.A.9: Labour force data, wage-employees: by gender

Ln(Real earnings)	(1) POLS	(2) FE	(3) FD	(4) POLS	(5) FE	(6) FD
Ln(Firm size)	0.112*** (0.0179)	0.0557** (0.0209)	0.0440* (0.0188)	0.125*** (0.0225)	0.0667 (0.0356)	0.0295 (0.0265)
Observations	922	922	426	490	490	174

Standard errors in parentheses

Specification as in Table 3.5, but split by men (columns 1-3) and women (columns 4-6)

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.A.10: Labour force data, self-employed: by gender

Ln(Real earnings)	(1) POLS	(2) FE	(3) FD	(4) POLS	(5) FE	(6) FD
Ln(No. Employees)	0.296*** (0.0801)	-0.0786 (0.123)	-0.0484 (0.120)	0.201* (0.0864)	0.181 (0.111)	0.166 (0.108)
Observations	655	655	303	1485	1485	758

Standard errors in parentheses

Specification as in Table 3.7, but split by men (columns 1-3) and women (columns 4-6)

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Chapter 4

Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession

Introduction

This chapter provides new evidence on the evolution of global interpersonal income inequality between 1988 and 2008 from a newly constructed dataset. This inequality concept measures inequality amongst all individuals in the world irrespective of their country of residency thus implicitly assuming a “cosmopolitan” social welfare function (Atkinson and Brandolini, 2010; Deaton, 2010) and translating the concern for within-country inequality to the global level. Over the period 1988-2008, the face of globalisation changed dramatically with the integration of many developing countries into the world economy. Not only did the world become more globalised but there was also rapid growth in some of the poorest countries. Global interpersonal inequality captures the effects of these shifts on both within- and between-country inequality, in contrast to studies using GDP per capita which ignore the former.

If, taking the long-run view, the Industrial Revolution had widened global inequality (Bourguignon and Morrisson, 2002), one might expect that the growth of China and

India would reduce inequality. Therefore, it might appear surprising that we find that inequality in the global income distribution, as measured for example by a Gini index, does not change very much over this period. However, this hides substantial re-ranking of country-decile-groups and changes in the regional composition of different parts of the global distribution. We first present new evidence on the evolution of global interpersonal inequality. We then look below the surface and analyse changes in the composition of the global distribution of income.

Measuring global inequality empirically is substantially more difficult compared with within-country inequality. In the absence of a global household survey, we need to resort to combining national surveys. Our database includes 565 household surveys across five benchmark years and each country-year observation is represented by the average income of ten income decile groups.¹ National surveys collect information in terms of local currencies, so we need to convert these into a common currency preferably adjusting for differences in the price level (across countries *and* over time).² We should point out that in constructing our global distribution we mix income and consumption surveys.³ We refer to them interchangeably, as is customary in this literature, although we are fully aware of the important differences between the two concepts.⁴

This chapter offers four main contributions to the study of global income inequality. First, we construct a new and improved database of national household surveys in response to criticism of earlier datasets (for example by Anand and Segal, 2008). The main source of data is the World Bank’s PovcalNet database which provides decile group average incomes for poor countries. The data for rich countries are mostly drawn from the Luxembourg Income Study (LIS). Second, this allows us to present more cred-

¹Each decile group is weighted by its population, so we measure interpersonal global income inequality where each individual is assigned the income of his or her income decile group.

²Analyses of national income inequality are often done on nominal incomes, thus assuming a common national price level. We follow these national studies by ignoring differences in the price level within countries (except for China, India and Indonesia where we allow for rural-urban price differentials) and across the income distribution. Also see the discussion in Chapter 1 of the thesis.

³The surveys we refer to as “consumption” surveys measure consumption expenditure. In the Becker (1965) framework, consumption is produced at home using both expenditure (supposedly including estimated expenditures on home-produced goods) and time.

⁴However, we improve on earlier approaches by keeping the type of survey (income or consumption) constant over time for a particular country.

ible results on global interpersonal income inequality between 1988 and 2008. Third, we create balanced and unbalanced panels of country-decile-groups for five benchmark years. This allows us to go further than the statements about which countries (and by how much) affected global inequality, by looking into a more disaggregated distribution of country-decile-groups. We can thus identify those country-decile-groups that have gained and lost most over this 20-year period. Fourth, we present one of the first comprehensive adjustments for missing top incomes in the study of global inequality.⁵ We offer valuable empirical insights because the effect on global inequality of non-response at the top within individual countries is unclear a priori (Deaton, 2005).⁶

We estimate the global Gini index to be around 0.70, which is towards the upper end of earlier estimates in the literature.⁷ It has fallen by approximately 2pp over this 20-year period, with the strongest decline between 2003 and 2008. For comparison, over the same period, the Gini index increased by 3pp in Germany and 6pp in the USA. Furthermore, 2pp is the difference in the Gini index between Germany and France in 2008. As another yardstick, generating such a decline in the Gini would require everyone to give up 2.29% of income to a common pool which is shared equally

⁵Pinkovskiy (2013) estimates nonparametric bounds on the global Atkinson index allowing for any country-level income distribution given fractile shares and a Gini index. With a sufficiently high non-response at the top, the direction of change of global *welfare* between 1970 and 2006 becomes ambiguous. He does not present bounds for global *inequality* measures which allow for non-response at the top.

⁶It might appear intuitive that stretching out the top tail would increase within-country inequality. However, it is important to recognise that non-response at the top also changes the mean and not only the variance (Deaton, 2005). The effect on global inequality depends on the extent of top income misreporting across countries, where in the global income distribution countries and their top income groups are, how far in the national income distribution non-participation begins (e.g. top 10% or top 5%), and how populous the countries and their top recipients are.

⁷Anand and Segal (2008) provide a comprehensive overview of the different estimates of global inequality. In Table 4.B.1 we provide an updated summary of the previous estimates drawing on a similar table in Anand and Segal. In Sections 4.3 and 4.4 we compare our results with these estimates in detail. Because studies differ in a number of aspects (e.g. use of per capita GDP), it is hard to identify the precise factor that explains why we obtain relatively high levels of inequality. One difference is that earlier studies used Geary-Khamis PPPs, which underestimate global inequality, an issue we will explain in detail in Section 4.1.3. Our methodology and results are most similar to Milanovic (2012). Bourguignon (2012) uses the same PPP exchange rates as our analysis but scales the survey mean to GDP per capita. He estimates a lower global Gini index in 1989 and the 20-year decline is approximately twice as fast (Table 4.B.1). This is similar to our estimates when we scale the survey mean to private consumption from national accounts (Section 4.4). This is consistent with a growing gap between national accounts and the household survey mean over this period. Warner et al. (2012) additionally assume a parametric within-country distribution and obtain very similar results to Bourguignon (2012).

(see p. 159). The time-series pattern is similar for alternative (relative) measures of inequality, such as the Theil indices. On the other hand, using an absolute Gini index, inequality in the global income distribution has increased by more than 20% between 1988 and 2008. Most of global inequality is accounted for by differences between countries, although this contribution has declined over time, suggesting that average incomes across countries have become more similar. The within-country component of global inequality has increased continuously over this 20-year period.

We present a number of robustness checks. The estimated change in the global Gini index is robust to using a balanced sample of countries over this period. However, as we will argue in more detail below, the changes in the global Gini over time are probably not significant with plausible standard errors.⁸ When we scale survey incomes to final household consumption from national accounts, we obtain a lower level of the global Gini index and a stronger decline over time.⁹ We also present a simple robustness check for underreporting of top incomes in household surveys. We treat the discrepancy between national accounts consumption and household surveys, an issue which has received considerable attention in the literature, as a proxy for the extent of “missing” top incomes. We obtain detailed top quantiles of the distribution by allocating this gap to the top decile group and fitting a Pareto distribution. The Gini index of this revised distribution is about 5pp higher and does not decline substantially between 1988 and 2008. For comparison, in 2008 the national Gini index is 5pp higher in Germany than in Sweden. Furthermore, reducing the global Gini from 0.759 (adjusted for missing top incomes) to 0.705 (baseline) in 2008 requires a global uniform tax of 7.08% (see p. 159). The difference in levels of around 5pp is primarily due to allocating the national accounts excess consumption to the top decile group rather than to the

⁸Pinkovskiy (2013) concludes that the direction of change in the global Atkinson index is ambiguous once he allows for flexibility in the country-level distribution linking the fractile shares and Gini index.

⁹We want to stress that scaling survey incomes to national accounts is not our preferred estimation strategy. We simply replicate the commonly adopted approach in this literature which scales survey incomes to GDP from national accounts. Ours is the first paper to compare the global Gini index based on household surveys with that scaled to national accounts consumption (rather than GDP), which should be preferred, as argued by Anand and Segal (2008) and as we explain in more detail below.

Pareto imputation itself, that is the elongation of the upper tail of the distribution.

The regional composition of the global distribution has changed substantially over this 20-year period. China has emerged from the bottom ranks of the global distribution, which has had a profound effect not only on the regional composition, but also on the overall shape of the global distribution. Both China's average income growth and change in income inequality have been exceptionally strong. India has grown more slowly, but its inequality has also grown much less. As a result of the growth in China (and, to a lesser extent, India), Sub-Saharan Africa now makes up most the bottom part of the global distribution.

Not surprisingly, China (particularly the urban part) is among the country-decile-groups which have grown most between 1988 and 2008. In Latin America, some country-deciles have also performed very well, but in contrast to China, these gains are mostly due to redistribution rather than growth. The new EU member states can be found both amongst the deciles that have grown fastest and slowest, similarly to Sub-Saharan Africa after 1993.

In constructing our global income distribution we have tried to be as careful as possible given the data constraints and corrected some of the biases in earlier studies.¹⁰ Nevertheless, sources of uncertainty over our estimates remain which are difficult (or impossible in some cases) to quantify.¹¹ We would suggest a conservative approach and conclude that the changes we observe over time are not statistically significant. We return to some of these issues in the conclusion and suggest ways in which we might

¹⁰For a comprehensive review, see Section 4.2. The improvements include (1) only using high-quality household surveys; (2) keeping the number of quantile groups constant across countries and over time; (3) splitting the three largest developing countries China, India and Indonesia into separate rural and urban regions, and accounting for differences in price-levels; (4) keeping the income or consumption concept constant over time for a given country; (5) using the most recent (and most reliable) set of PPP exchange rates and relying on domestic inflation rates for within-country comparisons. At the time of writing this chapter, the 2011 PPP exchange rates had not been published, so we use the 2005 PPPs and refer to them as the most recent round. The 2011 PPPs were released in April 2014. There is an ongoing debate about the reliability of these new exchange rates (Deaton and Aten, 2014; Ravallion, 2014a). Furthermore, our data end in 2008, so 2011 would be outside the period of investigation.

¹¹It is not clear how to derive standard errors for the global distribution which has been constructed from national household surveys each of which has its own sampling uncertainty. Furthermore, non-sampling uncertainty arising from measurement error in incomes and PPP exchange rates, mixing of income and consumption surveys, ignoring within-country inequality, and using unbalanced panels of countries might be more important (Anand and Segal, 2008).

address them in future work.

This chapter is structured in six main parts. First, we provide an overview of the literature on global inequality and poverty, the measurement of purchasing power parities (PPPs), the discrepancy between national accounts and household surveys and the underestimation of top incomes in household surveys, and outline our approach to all these issues. Section 4.2 summarises our data construction and methodology, including the welfare concept used and the Pareto imputation to account for top incomes. Section 4.3 provides summary statistics on our database, particularly the coverage of global and regional GDP and population, and presents the main results regarding inequality in the global distribution of income. We use different measures of inequality and investigate the changing regional composition as well as global growth incidence curves (GICs). In Section 4.4, we provide an upper bound on the global Gini by accounting for the underestimation of top incomes. Section 4.5 moves from a cross-sectional focus to a panel analysis, accounting for the movement of individual country-decile-groups in the global distribution. Section 4.6 presents conclusions.

4.1 Literature review and our approach

This chapter is related to a number of different strands of the literature. First, we summarise the literature on global inequality, and define what we mean by this term. Second, we consider the measurement of global poverty which has many issues in common with global inequality but in addition involves the choice of a global poverty line. Third, we explain the various problems associated with deriving PPP exchange rates and why we consider the rates used in this chapter to be the most robust. Fourth, we address previous work on (a) the underreporting of top incomes in household surveys, and (b) the discrepancy between national accounts measures of income or consumption and their household survey equivalents. We argue that the two issues are closely related and consider them jointly in our analysis.

4.1.1 Global inequality

Milanovic (2005) distinguishes between three concepts of global inequality. First, unweighted international inequality is the inequality in per capita incomes amongst the countries in the world. Second, population-weighted international inequality, or between-country inequality (Anand and Segal, 2008), measures inequality amongst persons by assigning everybody the per capita income of their place of residence. It thus ignores any within-country inequality. Third, global interpersonal inequality captures the inequality of individual incomes in the world by giving everybody their own income. Anand and Segal (2008) add concept zero, which is the inequality in total (rather than per capita) income among countries, i.e. not population-weighted.

In this chapter we focus on the third concept, global interpersonal income inequality, and whenever we use the term “global inequality” this is the concept we refer to. Unweighted international inequality (Concept 1) might be appropriate in studies of income convergence across countries (e.g. Barro and Sala-i Martin, 1992), but it is not a measure of interpersonal inequality, not least because the weight attached to individuals depends on their country of residence. Bourguignon (2011b) shows that between 1989 and 2006, unweighted international inequality (concept 1) has continued to increase, whereas global inequality (concept 3) has declined, as measured by the P90/P10 ratio.¹² This can be explained by the fact that some populous Asian countries have grown very fast, whereas some smaller (mostly African) countries lagged behind or even declined. Population-weighted international inequality (Concept 2) ignores within-country inequality which seems inappropriate given the widespread concerns on precisely this topic. It might be best seen as an intermediate (downward-biased) step towards global interpersonal inequality when survey data for measuring within-country distributions are not available (Milanovic, 2005).

Our interest in global interpersonal inequality is founded upon a concern for individual well-being and more specifically a concern for “global justice”, i.e. the distribution of resources amongst the world’s citizens which mirrors the concern for inequality at

¹²He uses GDP per capita combined with distributional data from PovcalNet and the OECD.

the country-level (Pogge, 2002; Singer, 2002). This cosmopolitan view of the world assumes a global social welfare function which treats persons the same irrespective of their country of residence (Atkinson and Brandolini, 2010). The cosmopolitan view is not shared by everyone and there exists a large philosophical literature on this issue (e.g. Nagel, 2005).¹³ However, it is important to recognise the increasing role played by international organisations. Not only are they the beginning of a world government, but also the global cosmopolitan view is the only one consistent with their constitutions. The cosmopolitan view is not the only reason for studying the global distribution. There might also exist an instrumental concern for global interpersonal inequality if extreme global inequality leads to increased international tension or to the rejection of western values. To address other concerns, a different global inequality concept may be more appropriate. For example, to understand unequal power in international politics, concept zero seems appropriate (Anand and Segal, 2008).

Changes in global inequality also capture some of the effects of globalisation (Anand and Segal, 2008). Over the 20-year period analysed in this chapter, the world economy has certainly become more integrated. Globalisation supported fast economic growth in many populous developing countries (mainly in Asia), thus reducing between-country inequality. At the same time, it is blamed for increasing inequalities within both rich (Autor et al., 2013) and poor (Goldberg and Pavcnik, 2007) countries, as the benefits are distributed unevenly across sectors, regions or factors, such as skilled and unskilled workers (Wood, 1995). Global interpersonal inequality captures both these effects in a unified framework. We want to stress that any estimate presented here cannot be given a causal interpretation since there exist no counterfactual world economies. In addition, 1988 was certainly not the start of globalisation. However, since then the pattern of global trade and capital flows has changed dramatically, with the integration of China (Haskel et al., 2012), other developing countries (Goldberg and Pavcnik, 2007) and Russia into the world economy.

¹³For example, Bhagwati (2004) calls the concern with global interpersonal inequality a “lunacy”, because individuals around the world “do not belong to a ‘society’ in which they compare themselves with the others” [p.67].

Ideally global inequality would be measured from a single representative global data source (such as a global household survey or administrative records), which would be analogous to measuring country-level inequality from national household surveys or administrative records. In the absence of such a data source, we have to rely on combining national household surveys. Most of the literature on global inequality uses distributional information (e.g. Gini indices or quantile shares) from secondary datasets, such as Deininger and Squire (1996), in combination with average incomes from national accounts (e.g. Bourguignon and Morrisson, 2002; Sala-i Martin, 2006; Pinkovskiy, 2013). Some studies also assume a parametric within-country distribution (e.g. Chotikapanich et al., 1997, 2012; Warner et al., 2012). Using average incomes from national accounts implies a rejection of the often available mean income or consumption from household surveys. According to Anand and Segal (2008), Milanovic (2002; 2005; 2012) is the only author who works directly with the household survey data without scaling to national accounts and we follow this approach in our baseline specification.

Anand and Segal (2008) offer a detailed overview of the literature on global interpersonal inequality to date. All studies agree that the level of inequality is high, with a Gini index between 0.63 and 0.71 in the 1990s.¹⁴ Because methodologies and data sources differ substantially (e.g. the use of national accounts aggregates, the estimation of within-country distributions, the use of different and often inconsistent PPP exchange rates), there exists substantial uncertainty over the direction of change in global inequality. Hence “there is insufficient evidence to reject the null hypothesis of no change in global interpersonal inequality over 1970–2000” [p. 91], to which Pinkovskiy (2013) agrees using a very different methodology (see footnote 5). In Section 4.3 and Table 4.B.1 we compare our estimates with these previous results.

Studies of global inequality typically use relative measures of inequality (with the exception of Del Río and Ruiz-Castillo, 2001).¹⁵ Using the Kolm (1976) definitions, a relative measure, such as the standard Gini index, is unchanged by a proportional

¹⁴The Gini index of the global *wealth* distribution is substantially greater at around 0.80 (Davies et al., 2011).

¹⁵Grün and Klasen (2008) use a measure of global welfare, which combines both inequality and average income.

transformation, i.e. it is scale-invariant. Absolute measures, on the other hand, are translation-invariant, which means that they are unchanged by an additive change. Atkinson and Brandolini (2010) argue that it is important to consider absolute inequality at the global level, because (a) income differences are so much larger at the global than the national level, and (b) absolute (e.g. global poverty) and relative measures of global inequality have evolved quite differently over time.¹⁶ Ravallion (2004, p.19) notes that the “perceptions that inequality is rising may well be based on absolute disparities in living standards”. Atkinson and Brandolini (2010) review a number of flexible parametric formulations which inter alia allow for different social weights applied to within-country and global inequality. In this chapter, we only estimate an absolute Gini index (defined in footnote 80), although we aim to explore these issues in more detail in future work.¹⁷ Absolute measures of inequality also naturally extend to the measurement of global poverty, an issue we turn to now.

4.1.2 Global poverty

The analysis of global inequality is closely related to the measurement of global poverty which has arguably received more attention.¹⁸ The two measures capture two aspects of the global distribution which have evolved quite differently over time. Whereas global inequality remains high with a Gini of around 0.70 (Table 4.3), global poverty has fallen substantially. Using a methodology similar to ours, Chen and Ravallion (2010b) find that the global poverty headcount rate has fallen from 52% in 1981 to 25% in 2005. Both require the same underlying data source, namely household surveys, and thus many of the same methodological issues arise, such as the use of PPPs discussed below. In addition, a measure of global poverty requires the choice of a poverty line. Chen and Ravallion (2010b) choose an international poverty line which is then converted

¹⁶Bosmans et al. (2013) find that population-weighted between-country inequality declined only when a strongly relative view of inequality is adopted.

¹⁷In their estimation using the Bourguignon and Morrisson (2002) data, Atkinson and Brandolini (2010) find that the time-series pattern of the absolute Gini index is similar to the more flexible measures.

¹⁸For example, the first of the Millennium Development Goals aims to halve global poverty between 1990 and 2015. Deaton (2010) argues that measures of global poverty are a strong determinant of the regions focused on by foreign donors and international development organisations.

into national currencies using PPP exchange rates and deflated over time using the domestic CPI. This approach is not without criticism and alternative approaches for setting poverty lines have been suggested (e.g. Srinivasan, 2001; Deaton, 2010; Reddy and Pogge, 2010).

Recently, (historic) global poverty rates were revised substantially upward. The new \$1.25 per day (in 2005 PPP) line increased the poverty headcount in developing countries in 2005 from 17% to 25% (compared with the old poverty line in 1993 PPP), adding 400 million poor people (Chen and Ravallion, 2010b). For China, this revision adds 130 million more poor people (Chen and Ravallion, 2010a). Deaton (2010) shows that the switch from 2005 to 1993 PPPs only had a small effect on the global poverty count. Instead, these revisions are mostly explained by changes in the global poverty line which is computed as an average of the poverty lines found in the poorest countries, which are changing over time. For example, India has become too rich to be included in this sample. On the one hand, the global poverty line should clearly be updated for new information about national poverty lines. Furthermore, it is reasonable to focus on poor countries because the line would cease to be an absolute line if the poverty lines of all countries, including the rich world, were included. But on the other hand, these substantial discontinuities in the poverty headcount are clearly troubling and maybe a weighted average of national poverty lines would be more appropriate (Deaton, 2010).

4.1.3 PPP exchange rates

Comparing incomes in different countries requires the use of exchange rates. If the law of one price held and there were no non-tradables, we could simply use market exchange rates (Deaton and Heston, 2010). This, however, is clearly not the case and market exchange rates would understate the real standard of living in poor countries, thus overstating global inequality (Anand and Segal, 2008).¹⁹ We use PPP exchange rates in order to account for price differences across countries. PPP exchange rates convert a given local currency into USD, the numeraire. Because we are dealing with

¹⁹ Alessandria and Kaboski (2011) explain the failure of the law of one price even for tradables by a model which assumes that poor country consumers have a comparative advantage in search activities.

household income or consumption, we use the PPP exchange rates for private consumption rather than the GDP conversion factors. Obtaining PPP exchange rates faces substantial conceptual (e.g. accounting for different consumption baskets) and practical (e.g. collecting comparable price data) challenges. Although we use the best available PPP exchange rates, the international price comparisons are a non-negligible source of error in our estimates of global inequality.

The first step in the computation of PPPs involves the collection of price data around the world coordinated by the International Comparison Program (ICP). The latest round of price comparisons refers to year 2005.²⁰ This round has the largest global coverage including 146 countries up from 117 in 1993, the previous round. China participated for the first time ever and India for the first time since 1985. In addition to the improvements in country coverage, the survey methodology was also improved in the latest round of the ICP, in particular in terms of product specifications.²¹

However, it is important to recognise that PPP calculations have become highly political given their role in measuring poverty (see footnote 18). Therefore, the source of any large revision needs to be assessed carefully. The issue of urban bias in price collection (that is, of price data collected disproportionately in urban areas) has received particular attention in the case of China, where the 2005 ICP round led to a substantial upward revision of the previous price level (which had been mostly based on guesswork). The price survey was conducted in eleven metropolitan areas which were chosen because they were likely to have outlets that sold the products compared in the ICP survey (Chen and Ravallion, 2010a). Chen and Ravallion (2010a) argue that the measured price level is representative of urban prices, but substantially overstates

²⁰At the time of writing, the 2011 PPP exchange rates had not been published (see footnote 10). Therefore we refer to 2005 as the most recent or latest round throughout.

²¹In previous rounds lower quality goods in poor countries were compared with goods of higher quality in rich countries, and as a result price levels in poor countries understated (Deaton and Heston, 2010). The product specification has to find a balance between (1) being accurate by matching the same products across countries and (2) avoiding pricing rarely consumed and thus non-representative and expensive goods. The well-known Big Mac Index is an extreme attempt to price the exactly same product, which is a relatively luxurious good in poor countries, and as a result overstates the price level compared with PPP exchange rates (Ashenfelter, 2012). Deaton (2010) argues that the latest round of PPPs has gone too far towards choosing precisely matched, but unrepresentative goods in poor countries, thus overstating price levels in these countries and thus increasing global inequality.

the rural price level. In this chapter we follow the approach adopted by Chen and Ravallion (2010b) of treating the official PPP rate as representative of urban China and using a downwardly adjusted PPP rate for rural China.

Errors in the price collection stage also arise from standard sampling error in price surveys (Anand and Segal, 2008) and difficulties in pricing hard to compare products such as housing (Deaton and Heston, 2010). But despite all these issues “[t]here are many fewer health warnings for the 2005 round than for earlier rounds” [p.2] and “[i]n many cases, the prices collected under the auspices of the ICP may be more accurate than the GDP numbers with which they are combined” [p.4] (Deaton and Heston, 2010).

The second step in the estimation of PPP exchange rates involves the computation of a price index, i.e. a particular weighting scheme which combines the national prices collected in the first step (our discussion here draws heavily on Deaton and Heston, 2010). In the most recent 2005 round, the ICP has used the index due to Eltetö and Köves (1964) and Szulc (1964) (EKS). The Penn World Tables and earlier estimates by the World Bank used the Geary-Khamis (GK) method (Khamis, 1972). The EKS index is a multilateral price index which combines all the bilateral Fisher price indices.²² More precisely it is the geometric mean of all the indirect Fisher indices between the base country and the country of interest.²³ The EKS index satisfies transitivity, so we have one index per country instead of a matrix of indices. But the EKS method violates the “independence of irrelevant country” property, because the index between any two countries is affected by prices and weights in third countries (precisely because it combines the indirect Fisher indices).

The GK index compares domestic prices with world prices. The problem is that in the computation of these world prices, the weight attached to a particular country

²²The Fisher index is the geometric mean of the Paasche and Laspeyres indices. Its weights thus take into account both the reference and comparison countries. Importantly, it naturally obeys the reversal property, which means that the price level of country 1 based on country 2 is the inverse of the price level of country 2 based on country 1.

²³Suppose we are interested in estimating Ghana’s price level (relative to the US, the base country). There exists a Fisher index which goes from the US to Germany and a separate Fisher index which goes from Germany to Ghana. The EKS index is a geometric mean of all these indirect Fisher indices.

depends on its physical volume of consumption. Thus in practice, the rich countries dominate these composite world prices, which are most similar to Hungarian prices (Nuxoll, 1994). Since goods that are relatively expensive in rich countries (say, services) tend to be consumed in relatively large quantities in poor countries, precisely because they are cheaper, the use of a GK index tends to overestimate the value of consumption in poor countries.²⁴ This is a manifestation of the Gerschenkron (1947) effect (or substitution bias) which says that a country’s consumption is overvalued when evaluated at the prices of another country, and the further apart the two price vectors, the greater the overvaluation. The Fisher index, and thus the EKS index does not suffer from this bias because it averages the consumption weights from both countries, making “a compromise that is arguably the best that can be done in the circumstances” (p.11, Deaton and Heston, 2010).²⁵ As a result of the Gerschenkron effect, the GK index understates global inequality (and global poverty, see Ackland et al., 2013). Deaton and Heston (2010) compute population-weighted between-country inequality of GDP per capita using different indices. They obtain a Gini index of 0.533 for EKS and a value of 0.527 for GK, which is a small difference. Hence, the choice of index method is likely to be of secondary importance. Put differently, most of the effect of switching from 1993 to 2005 PPPs seems to come from the new price data.

In summary, we use the EKS index as suggested by a number of authors (Anand and Segal, 2008; Deaton and Heston, 2010; Ravallion, 2010b; Ackland et al., 2013). For other purposes, the GK index which satisfies additivity might be more relevant.²⁶ Other methods, such as the Afriat indices suggested by Dowrick and Akmal (2005), lack empirical applicability (Anand and Segal, 2008). We use a single PPP exchange rate per country (differentiating between rural/urban China, India and Indonesia), thus

²⁴The denominator of the GK index uses (approximately) Hungarian prices weighted at (say) Ghanaian budget shares, which depresses the PPP exchange rate and thus overvalues Ghanaian consumption.

²⁵For example, consider alcohol in Bangladesh, where it has a small share, but a relatively high price (see Deaton and Heston, 2010). The Fisher index strikes a compromise between two extreme positions: Bangladeshi budget shares would understate Bangladeshi prices, whereas OECD shares would overstate them.

²⁶For example, this might be important when studying the composition of GDP because, in contrast to the GK index, components converted at EKS conversion factors are not additive.

ignoring any consumption differences along the income distribution, which have been discussed in the context of global poverty measurement.²⁷

The final issues in the use of PPPs are how to extend them over time and, related to this, how to deal with revisions of the PPP price indices. Similar to Chen and Ravallion (2010b), we compare prices across countries only once using the most recent ICP round, relying on domestic consumer price inflation for the within-country comparisons. In other words, our approach only requires one ICP round (and with every new round of PPP exchange rates we would simply start anew): domestic local currency units in any year are converted into domestic 2005 prices using a domestic CPI deflator, and to these (constant-price) local currency units we then apply the 2005 PPP exchange rate obtained from direct price comparisons by the ICP.²⁸

Conceptually, our approach is simple because it keeps the comparisons over space and over time separate. All our within-country comparisons are independent of international prices and only depend on domestic prices, which is attractive not least because domestic prices are appropriate for assessing the trade-offs at the country-level (Nuxoll, 1994). We avoid errors arising from mixing different ICP rounds (Anand and Segal, 2008) and we only use the most recent and arguably most robust round of PPP rates. Although not without error, domestic inflation rates are the best data available for making within-country comparisons over time (Ravallion, 2010a).

However, our approach has been criticised for two reasons. First, PPP exchange rates and inflation rates measure different things (Reddy and Pogge, 2010).²⁹ Different rounds of PPP comparisons reflect the quantities and prices in that year across the world (Deaton and Heston, 2010). On the other hand, the CPI tries to account for

²⁷Reddy and Pogge (2010) argue that consumption PPPs should not be used in the measurement of global poverty because the consumption baskets of the poor are systematically different from the rest of the population. For example, the poor might face different prices because of where they buy (Frankel and Gould, 2001) or how much they buy (Rao, 2000). Also see the discussion in Chapter 1 of the thesis. However, Deaton and Dupriez (2011a) find that re-weighting commodity baskets to account for different consumption baskets of the poor and non-poor has very little effect on the estimated PPP exchange rates, which is also reassuring for our estimation strategy.

²⁸In contrast, the Penn World Tables use an averaging procedure to fill the intervening years between consecutive ICP rounds (more details in Deaton and Heston, 2010).

²⁹In addition, the weights in the PPP exchange rates are derived from national accounts, whereas the CPI typically uses weights from consumption surveys.

changes in the price level within a country over time. Hence going further back in time, we would lose international comparability. We argue in our defence that (1) we are extending the index backwards by at most 20 years, i.e. not too far backwards, and (2) given that the previous rounds of the ICP were of questionable quality, our approach appears the most robust.

Second, whenever there is a new ICP round, our approach would simply re-compute all the historical estimates anchoring to that new base year. These revisions can be substantial as in the case of switching from the 1993 to the 2005 base year. Milanovic (2012) revises the Gini index of global income inequality upward by 5pp, which is a substantial change (Deaton, 2010).³⁰ The shrinking of the Chinese economy by some 40%, the addition of some 130 million poor people, and the changing regional distribution of poverty has probably received most attention (Reddy and Pogge, 2010). It is true that any choice of base year is essentially arbitrary (Reddy and Pogge, 2010) and “if revisions can have such enormous effects, it is difficult to place much confidence in any set of poverty estimates” (p. 128, Deaton, 2001). However, we offer three arguments in favour of our approach. First, as discussed above, revisions of the global poverty count arose mostly as the result of changes in the poverty line, instead of the change in the base year (Deaton, 2010). Second, it could be argued that the revisions are the result of the “index number problem” of making international welfare comparisons and thus cannot be resolved even in principle (Deaton, 2001). Third, we have argued above, and many authors agree³¹, that the most recent ICP round constitutes an improvement both in terms of method (EKS index vs. GK index) and country coverage (e.g. the inclusion of China as a benchmark country). Switching from a GK to an EKS index would increase global poverty by itself (Ackland et al., 2013). Furthermore, increasing productivity in the tradable goods sector puts upward

³⁰Warner et al. (2012) also obtain an upward revision of the Gini index by 5pp when they update the estimates by Chotikapanich et al. (2012) to the 2005 PPPs (Table 4.B.1), although they use a rather different methodology (using GDP per capita and parametric within-country distributions).

³¹According to Deaton and Heston (2010) “many of the most marked differences [between the 2005 round and earlier rounds] reflect corrections of problems in earlier rounds” [p. 2]. And Deaton (2010) writes that “[g]iven the clear improvement in the successive ICPs, I would certainly not advocate the retention of outdated and flawed data” [p. 9].

pressure on the overall price level, which can explain two-thirds of the increase in the Chinese price level index from 1993 to 2005 (Ravallion, 2010b).³² The upward revisions of the African and Asian-Pacific PPPs are behind the increase in the estimated global inequality (Deaton, 2010).

4.1.4 National accounts and household surveys

Typically, per capita household consumption in national accounts exceeds the average consumption or income recorded in a household survey (Deaton, 2005).³³ Moreover, as Deaton shows, the discrepancy appears to have increased over time not only in India (which is somewhat of a *cause célèbre* in that respect) but also in rich countries like the United States and Great Britain. Studies of global inequality differ in their view of how to account for this discrepancy. In our main specification, we follow the approach suggested by Anand and Segal (2008), and simply use the average income recorded in the survey.

As mentioned before, other papers (except Milanovic, 2002; 2005; 2012) have anchored the income level to the national accounts and combined it with distributional information from household surveys: for example, Sala-i Martin (2006), Bhalla (2002), Bourguignon and Morrisson (2002), Bourguignon (2012), Chotikapanich et al. (2012) and Warner et al. (2012) set a country's income level equal to GDP per capita. Anand and Segal (2008) argue that GDP per capita "is not a suitable measure of household income" [p. 67] and should not be used to anchor household survey means, since it includes items such as depreciation, retained corporate earnings or taxes which are not distributed back to households, all of which are only remotely related to household income. Similarly, Atkinson (2013a) proposes a new headline indicator for assessing economic performance, which is a more direct measure of household living standards

³²The price level index is defined as the PPP exchange rate divided by the market exchange rate. For China, it rose from 0.25 in 1993 to 0.52 in 2005 (Chen and Ravallion, 2010a).

³³In some - mostly African - countries, per capita national accounts consumption is actually lower than that found in household surveys. Deaton (2005) argues that this might be explained by underestimation in national accounts rather than by problems in household surveys. Devarajan (2013) discusses the measurement problems that plague GDP, population and poverty statistics in Africa.

than GDP.³⁴ Furthermore, there exists “a basic incongruity in assuming that the relative within-country distributions are measured acceptably well by surveys but their means are not” (Anand and Segal, 2008, p. 70). In addition, the replacement of the survey mean by a typically larger GDP per capita implies an equi-proportional adjustment of all incomes. Such a “proportional” adjustment, which scales up incomes by the same relative amount across the entire distribution, is unlikely to be correct if the national accounts gap is related to the underestimation of top incomes in surveys, as we will argue in Section 4.1.6.

Compared with GDP, final household consumption expenditure from national accounts is closer to household income (or consumption) recorded in surveys (Anand and Segal, 2008).³⁵ However, it must be noted, that the data and methods used to estimate national accounts consumption are not necessarily more reliable than household surveys (Anand and Segal, 2008; Deaton, 2005).³⁶ There are also definitional differences between household surveys and national accounts consumption, such as the inclusion of the imputed value of owner-occupied housing, imputed financial services, and consumption by non-profit institutions (Deaton, 2005).

4.1.5 Top incomes

A recent literature using tax records argues that top incomes are understated in standard household surveys. This literature studies inequality at the very top of the distribution, typically expressed in top income shares, e.g. the share of total income received by the top 1% (Atkinson and Piketty 2007; 2010). Tax data might provide more accurate information on top incomes for a number of reasons: First, it might be harder

³⁴Atkinson also suggests including distributional characteristics rather than per capita GDP, which we do by looking at the ten decile groups. His argument is closely related to the Stiglitz-Sen-Fitoussi commission, which concludes that “trends in citizens’ material living standards are better followed through measures of household income and consumption” (p. 13, Stiglitz et al., 2009).

³⁵It might be argued that *income* recorded in household surveys should be approximated with GDP (rather than national accounts consumption). However, Anand and Segal (2008) argue that even in this case, national accounts consumption is to be preferred because of the unrelated components included in GDP, as mentioned above.

³⁶The measurement issues in national accounts include the measurement of illegal transactions or the measurement of intermediate inputs. Furthermore, because consumption is computed as a residual, measurement errors are compounded (Anand and Segal, 2008).

to enter the gated communities of the rich than to conduct surveys in poor areas, so survey non-response would increase with income (Groves and Couper, 1998). Second, the top 1% are rare by definition, so a household survey with a standard sample size of a few thousand would offer top share estimates with a low precision, or might miss these people altogether. On the other hand, the tax data intentionally oversample the rich. Third, aspects of survey design such as top-coding or the elimination of “outliers”, manipulate top incomes. However, tax data are not without problems, e.g. due to tax evasion and avoidance which may be particularly strong in developing countries.

There is some evidence to support the argument that top incomes are missing in household surveys. Alvaredo (2010) finds that a household survey in Argentina records no observations with incomes exceeding \$1 million whereas the Argentinian tax data contain close to 700 observations in that range. In a comparison of household surveys from 16 Latin American countries, Székely and Hilgert (1999) find that the 10 richest households in the survey receive incomes similar to a managerial wage. It would appear plausible that the top capital owners in these countries receive substantially greater incomes than a manager. Some studies compare top shares estimated from household surveys and tax data, and in some cases obtain similar results, although this typically depends on the availability of exceptional surveys which have sufficient sample sizes and are not subject to top coding (Burkhauser et al. (2012) using internal CPS data, Leigh and van der Eng (2009) for Indonesia and Morival (2011) for South Africa).

Given that tax data appear to be more accurate at measuring top incomes and household surveys offer more precise information about the rest of the distribution, a natural next step would be to combine the two sources of information to obtain a complete distribution of income. However, the still sparse availability of tax data across countries limits the usefulness of such an exercise for the purpose of analysing the global distribution. In addition, the population and the welfare measure are fundamentally different between the two data sources which makes such an exercise difficult.³⁷

³⁷This is discussed in more detail in Chapter 2 of this thesis. In the tax data, the unit of analysis is a tax unit, which depending on the jurisdiction could be a married couple or an individual. A household would typically be bigger than a tax unit. The tax data literature uses taxable (and usually before-tax) income, whereas the household survey collects disposable income. Taxable income excludes

4.1.6 Missing top incomes and the national accounts discrepancy

The underreporting of top incomes in household surveys and their discrepancy with national accounts are closely connected issues. It is reasonable to expect, and there is some empirical evidence to corroborate it, that the discrepancy between surveys and national accounts is not distribution-neutral and could at least partially be due to non-participation of the rich in household surveys (Mistiaen and Ravallion, 2003; Korinek et al., 2006).³⁸ Deaton (2005) points out that because national accounts consumption tracks money rather than people, national accounts data are more likely to capture large transactions. Using Indian tax record data, Banerjee and Piketty (2010) find that a significant part of the discrepancy between consumption growth in national accounts and household surveys can be accounted for by underreporting of the rich. Finally, it could be argued that household surveys offer a good approximation to the bottom 90% of the distribution (thus, however, ignoring any underreporting of incomes among the very poor).³⁹

In Section 4.4 we allocate the gap between household final consumption in national accounts and household surveys to the top 10% of the distribution and obtain more disaggregated top quantiles by fitting a Pareto distribution to the upper tail.⁴⁰ Our

some real income, such as tax-exempt interest on government bonds. On the other hand, household surveys typically measure capital incomes and gains poorly compared with tax data, which cover at least taxable capital incomes and gains. Because the tax data typically do not contain sufficient information to construct units and incomes which are similar to those in a standard household survey, the only possibility is to construct tax units and taxable income in the household survey, as in Alvaredo (2010). In the final step of such an exercise, one needs to assume which parts of the true distribution are represented by the tax data and which are covered by the household survey. CBO (2012) matches US tax records with household survey records using income, and adds non-taxable income (e.g. transfers or in-kind income) from the survey. Armour et al. (2013) also match records by income but they add capital incomes to the household survey because these incomes tend to be poorly measured in these data.

³⁸Anand et al. (2010) argue that increasing income shares of the rich can explain why the discrepancy is increasing in countries such as China or India.

³⁹The coverage of the poor may be incomplete because surveys may exclude the homeless and institutionalised populations (see Carr-Hill, 2013), or because of sampling issues (excluding remote and probably poorer regions). Thus the very bottom of the distribution may be truncated. But, in addition, incomes of the poor included in the survey may be mismeasured due to extensive home production or benefits received in kind which may not always be included.

⁴⁰The Pareto distribution is known to fit the upper tail of income distributions well (Atkinson et al., 2011).

approach builds on Atkinson (2007b) who uses a Pareto imputation in combination with the Bourguignon and Morrisson (2002) data. Atkinson uses GDP per capita to anchor the country mean, thus spreading the discrepancy between national accounts and household surveys *evenly* across the distribution. For the very top he “elongates” the distribution by using a Pareto imputation.⁴¹ We call Atkinson’s approach the “proportional adjustment with Pareto tail”. By contrast, our methodology proposes to allocate the “excess” consumption recorded in national accounts only to the top decile group and to use a Pareto imputation, thus both increasing the mean and changing inequality. We call this adjustment “top-heavy with Pareto tail.”

We justify the adjustment of the survey mean to national accounts consumption on the grounds of missing top incomes. This is however open to criticism. It might be argued that some elements of national accounts consumption should (1) be excluded altogether, or (2) spread more widely along the distribution than the top 10%. For example, some of the discrepancy between national accounts consumption and household surveys is related to differences in definition, such as the inclusion of expenditures by non-profit institutions serving households (NPISH) or imputed consumption of collectively provided goods.⁴² We ought to subtract these components from national accounts consumption but sufficiently detailed data are not available separately for a large number of countries. Other sources of the discrepancy, such as imputed rents of owner-occupiers in the national accounts, could be included (if they are not estimated by household surveys), but should perhaps be spread further along the distribution than just the top 10%.⁴³ For these reasons our estimates should be seen as an approximate first step, which is likely to be an upper bound, in the absence of a more careful analysis using unit-record data.⁴⁴

⁴¹The Pareto imputation does not add “new” observations, but rather “stretches out” the top decile group. This implies that the imputation by itself does not change the mean.

⁴²According to Deaton (2005), NPISH consumption accounted for 3.9% of total consumption in the UK in 2001, up from 2.1% in 1970. In addition, he argues that this share might be even higher in poor countries, although there exist no data.

⁴³Deaton (2001) argues that in India approximately half of the gap between national accounts consumption and household surveys is due to imputed rents.

⁴⁴In order to address the undersampling and underreporting issues in surveys, Anand and Segal (2008) propose to estimate a parametric within-country distribution with a plausible upper tail from unit-record data. This could be done with standard models for truncated data and would address the

4.2 Data construction and methodology

4.2.1 Data source

The data used in this chapter consist of country-year decile group average income or consumption covering the period 1988 to 2008. The data come from a number of sources. PovcalNet is the starting point of our database, contributing more than two-thirds of the surveys.⁴⁵ PovcalNet is the compilation of a large number of household surveys by the World Bank research department. It has been mostly used to estimate global poverty, as in Chen and Ravallion (2010b), and thus lacks data on rich countries. From PovcalNet we obtain average per capita incomes, already converted in 2005 \$PPP, and decile group income shares, which we combine to compute decile group average incomes.⁴⁶

Next, we merge the updated World Income Distribution (WYD) data (Milanovic, 2012). We convert these data into country-year deciles in order to obtain a consistent database.⁴⁷ PovcalNet and WYD provide almost 98% of all data. Where possible we fill remaining gaps with data from LIS, the British Household Panel Survey (BHPS), the European Union Survey of Income and Living Conditions (SILC) and data from country statistical offices.⁴⁸ Overall we end up with 565 surveys across the five benchmark years

mismeasurement in both the mean (where we used national accounts consumption) and the dispersion (where we used the Pareto distribution).

⁴⁵PovcalNet: the online tool for poverty measurement developed by the Development Research Group of the World Bank, <http://iresearch.worldbank.org/PovcalNet>. Data downloaded on 29 July 2012, which refers to the last data update on 28 February 2012.

⁴⁶PovcalNet uses grouped data derived from household surveys to calculate these decile group shares. They are estimated from Lorenz curves fitted to the population-weighted (accounting for household size and sampling weights) distribution of per capita household income or consumption. Both Generalised Quadratic and Beta Lorenz curves are estimated and the functional form with the better fit is chosen.

⁴⁷The vast majority of country-year observations are already in deciles or in equally spaced quantiles (e.g. ventiles), so they can be easily converted. In total 13 country-years were imputed by fitting a log-normal Lorenz curve using the “ungroup” command included in the DASP Package (Abdelkrim and Duclos, 2007). This procedure implements the Shorrocks and Wan (2008) adjustment thus ensuring that the fitted Lorenz curve matches the original points. We chose a log-normal functional form and fitted the Lorenz curve on 2000 observations, as suggested by Shorrocks and Wan (2008). Minoiu and Reddy (2012) show that for estimating the global income distribution it is better to fit a parametric Lorenz curve rather than a kernel density.

⁴⁸The sources for the final database (country-years in parentheses): PovcalNet (379), WYD (173), LIS (8), SILC (2), and one survey each from BHPS (Bardasi et al., 2012), Statistics Finland and Statistics Portugal. The BHPS data were accessed via the UK Data Service.

1988, 1993, 1998, 2003 and 2008 (Table 4.1). The vast majority of surveys cover the entire country except for several, mostly Latin American, countries which survey only urban areas.⁴⁹ We treat these surveys as representative of the entire country.⁵⁰

Each country's distribution is represented by the average incomes of the ten decile groups. This is not dissimilar from other studies in the literature such as Bourguignon and Morrisson (2002) who use 11 quantiles. As argued by Anand and Segal (2008) this ignores any within-decile inequality, thus understating within-country inequality and global inequality. Our choice of decile groups was dictated by PovcalNet, where more detailed information is unavailable.⁵¹ Furthermore, because of consistency and ease of discussion, we decided to use deciles also in those surveys where more detailed information was available. Evidence from a different version of the data where we have percentiles, suggests that the effect on the global Gini is around 0.5pp in 2008.⁵²

4.2.2 Survey selection and the definition of regions

The surveys included in the database need to meet two conditions: First, they need to be within two years of a benchmark year (as defined below).⁵³ Second, they need to be at least three and no more than seven years from the previous and next survey. The rationale of the second condition is to avoid surveys that are too close or too far apart from the interval of five years since some of the analysis is based on the assumption that five-year intervals hold throughout the sample. Table 4.1 shows how closely the

⁴⁹These countries include Argentina, Colombia, Ecuador, Honduras, Micronesia and Uruguay, and are all taken from PovcalNet.

⁵⁰This is likely to understate within-country inequality in these countries, because we expect the rural areas to be poorer compared with the urban areas. However, the share of the rural population in several of them (e.g. Argentina, Uruguay) is small.

⁵¹Using PovcalNet, Segal (2011) obtains more detailed fractiles (limited to represent at most 5 million people) by accessing the detailed estimation code of the Lorenz curves which PovcalNet fits on the grouped data. This is cumbersome and was not feasible given the large number of surveys we deal with, but may be addressed in future work. Furthermore, it is important to bear in mind that these more detailed fractiles require a parametric assumption on the Lorenz curve.

⁵²Those data have more limited global coverage and we do not separate by rural and urban regions for China, India and Indonesia. In this dataset, moving from percentile to decile groups reduces the global Gini from 68.9% to 68.5% in 2008.

⁵³Our approach differs from Chen and Ravallion (2010b), who use an interpolation method to fill missing years.

survey and benchmark years are aligned.⁵⁴ The choice of benchmark years is essentially arbitrary and we followed Milanovic (2012) in choosing 1988, 1993 and 1998. Global household survey coverage is very poor prior to 1988. 2003 and 2008 were chosen in order to obtain equally spaced benchmark years. Compared with Milanovic (2012), we managed to obtain a closer fit to the benchmark year in all years with roughly 3/4 of surveys conducted within one year of the benchmark year.

We use a mix of income and consumption surveys, as is customary in this literature. Although there are obviously important differences between income and consumption, we refer to them interchangeably.⁵⁵ We do not adjust for differences between income and consumption surveys because any such adjustment, applied to deciles, would be arbitrary.⁵⁶ One of the innovations of our database is that we restrict the welfare measure to be the same over time for a given country. This avoids any spurious changes arising from a change in the welfare concept being used.⁵⁷ For each country, income or consumption was chosen so as to maximise the number of benchmark years covered (subject to the two conditions in the previous paragraph).⁵⁸ As Table 4.1 shows, in the overall sample the number of consumption and income surveys is almost equal. In earlier years, the majority of surveys collected information on income, whereas in recent years the reverse is true. This can be explained by the improved survey coverage

⁵⁴For example, for the benchmark year 2008, 7.4% of surveys were conducted in the year 2006. In the case of PovcalNet, the survey year is sometimes not an integer, which would happen if a survey was conducted over more than one year. In all cases we took the start year of the survey, i.e. we rounded down the non-integer years.

⁵⁵Inequality tends to be lower in terms of consumption than in terms of income (Deininger and Squire, 1996), and (at least in the USA) has increased by less in consumption-terms (Fisher et al., 2013). In the full PovcalNet data (not the sample used in this chapter), the average Gini index of consumption surveys (37.98) is approximately 10pp lower than the average Gini over the income surveys (48.44). This is more than the Gini adjustment proposed by Li et al. (1998) of 6.6pp.

⁵⁶For example, the income data could be scaled down by savings in the national accounts (i.e. the gap between national consumption and national income) (Deaton, 2001). We did not adopt such an approach because of (1) the issues with national accounts data discussed above and (2) evidence that the savings rate is not invariant with income (Dynan et al., 2004). Chen and Ravallion incorporated such an adjustment in early estimates of global poverty but abandoned it in later work (Chen and Ravallion, 2004).

⁵⁷Income and consumption inequality might move in different directions over the same period of time (e.g. Krueger and Perri, 2006).

⁵⁸For Bulgaria, Botswana and Croatia the number of benchmark years with consumption and income information were the same. In all cases we chose the type of survey in order to maximise the number of surveys drawn from PovcalNet. For Nicaragua, we had both types of surveys in all years from PovcalNet. We chose income surveys since this is the prevalent type of survey in Latin America.

of poor countries where consumption surveys are more common (with the exception of Latin America).⁵⁹

We have grouped countries into eight regions. The first group consists of “mature economies”, which are the EU-27 countries plus the high-income countries in the world.⁶⁰ We treat India and China as regions in their own right. The remaining groups are defined as residuals according to the geographic regions used in the WDI.

4.2.3 Welfare concept

We are interested in analysing the global distribution of (annual) per capita income (in 2005 PPP adjusted USD). Our unit of analysis is the individual and the reference unit for measuring incomes is the household. Per capita incomes ignore any economies of scale in household consumption and within-household inequality, which is the approach followed in PovcalNet, our main data source.⁶¹ Per capita incomes have the advantage that they are simple to compute and have natural counterparts in the national accounts (which do not compute equivalised incomes). The effect of using a different equivalence scale on world inequality is not clear a priori.⁶²

In our database, each country-year distribution of per capita income is represented by the average incomes of the ten decile groups.⁶³ In the analysis, each decile group

⁵⁹In China, India and Middle East/North Africa, our database only uses consumption surveys (with the exception of All-China where income surveys are used). In Sub-Saharan Africa 98% of surveys are consumption surveys. In other Asia, 91% are consumption surveys. On the other hand, in the mature economies and Latin America, 97% and 96% respectively are income surveys.

⁶⁰To be precise, the mature economies include EU-27, Australia, Bermuda, Canada, Hong Kong, Iceland, Israel, Japan, Korea, New Zealand, Norway, Singapore, Switzerland, Taiwan and USA. Maddison (2008) uses “Western Europe and its offshoots”. Milanovic (2005) uses “Western Europe, North America and (the rich) Oceania - WENAO”. Essentially, we expand these two classifications by the rich Asian countries.

⁶¹Chen and Ravallion (2010b) justify their approach by arguing that for poor people there is little scope for economies of scale.

⁶²In their study of the LIS data, Atkinson et al. (1995) find that the inequality of per capita household income is greater than the inequality of household income adjusted by a square root equivalence scale. The precise effect on cross-country comparisons of inequality depends on the joint distribution of family size and income.

⁶³For Switzerland in 2008, the bottom decile had average income of zero. We have recoded this to missing, because otherwise the bottom 10% in Switzerland would be the poorest people in the world. Hence we have one missing observation in 2008. In general, when working with the micro data, we dropped observations with zero income, which corresponds to the convention used by LIS in its key figures. While being clearly not fully satisfactory, this procedure was adopted because it is not possible to distinguish between missing values and zeros in some datasets (Ravallion, 2014b).

is weighted by its population (i.e. 10% of the national population from the World Development Indicators (WDI)). Whenever we are interested in the performance of our estimation or database, e.g. the split between consumption and income surveys or the size of the Pareto constants, observations are unweighted, as pointed out in the tables.

We use consumption PPPs to account for price differences across countries. Incomes obtained from PovcalNet are already converted to 2005 PPP USD. For the additional surveys we replicate the approach in PovcalNet: as explained before, and after accounting for currency conversions⁶⁴, we convert incomes into local currency units in 2005 prices using domestic consumer price indices (CPIs).⁶⁵ Then, we apply the 2005 PPP consumption exchange rates to convert into international dollars.⁶⁶ It is important to note that PPP exchange rates only exist at the country-level, so we ignore any price differences which exist within countries. As a result we probably overstate within-country inequality. As mentioned before, we treat the rural and urban areas of the three most populous developing countries China, India and Indonesia as separate “countries”. Due to a lack of disaggregated data, we assume a common CPI for rural and urban areas, but allow for different PPP exchange rates.⁶⁷

4.2.4 Pareto imputation and scaling to national accounts consumption

We allocate the excess of national accounts consumption over household surveys in two steps. First, we adjust the country mean to equal the maximum of the survey mean and national accounts consumption.⁶⁸ Second, we re-compute the decile group shares

⁶⁴Currency conversions include changes in the currency being used, such as the formation of the Eurozone, and currency redenominations as often observed in high-inflation environments.

⁶⁵As a first step, we use domestic CPI figures from the WDI. Where these are unavailable, we resort to the IMF’s World Economic Outlook or the country’s statistical offices directly.

⁶⁶The WDI does not provide PPPs for Kosovo, so we used the implied PPP GDP conversion factors reported in the IMF’s World Economic Outlook.

⁶⁷Data for China and Indonesia are provided by PovcalNet so they already incorporate adjustments for differential costs of living in rural and urban areas. For India, where we added one survey from the WYD, we use urban and rural PPP exchange rates of Rs 17.24 and Rs 11.40 respectively which are given in Ravallion (2008) and used by PovcalNet.

⁶⁸To be precise, the revised mean is thus equal to $\max(\text{national accounts consumption}, \text{survey mean})$. In the majority of cases, national accounts consumption exceeds the income recorded in the survey,

for all decile groups except the top using the original decile group average incomes and the adjusted mean. The share in total income of those decile groups therefore decreases. Third, we compute the new top 10% share as the difference between 100% and the sum of the revised shares of the bottom nine decile groups. We use the revised top 10% and top 20% shares in the Pareto imputation.

We obtain household final consumption expenditure (in 2005 PPP USD) from the WDI for the survey years.⁶⁹ It is important to note that the sample changes in this part of the analysis for two reasons. First, due to the unavailability of macro data, we lose some country-year observations. Second, because of a lack of disaggregated macro data, we use the country-wide distributions in the case of China, India and Indonesia (where we previously used separate rural/urban distributions). In the case of China, we now use income surveys where we previously used consumption surveys.⁷⁰

We use a Pareto imputation to split the top decile group into smaller quantiles. We choose to split the top 10% into P90-P95, P95-P99 and P99-P100. The resulting dataset thus consists of 12 (uneven) fractile groups per country (which are weighted by the appropriate population in the analysis). The implicit assumption is that the top decile of our data follows a continuous Pareto distribution. Let H_i be the cumulative population share of individuals with incomes greater or equal to y_i , e.g. the top 10%. Let S_i be the share of total income received by this group. Atkinson (2007b) shows that for the Pareto distribution, the relative share of two top groups is given by

$$\log \left(\frac{S_i}{S_j} \right) = \left(\frac{a-1}{a} \right) \log \left(\frac{H_i}{H_j} \right) \quad (4.1)$$

so the revised mean equals national accounts consumption. When the survey mean is greater than per capita national accounts consumption, we keep the survey mean. Hence we implicitly assume that household surveys can only underestimate and not overestimate incomes.

⁶⁹It is defined as “the market value of all goods and services [...] purchased by households” (<http://data.worldbank.org/indicator/NE.CON.PRVT.PP.KD>). It includes durable products, imputed rent for owner-occupiers, payments to obtain permits and licenses, and expenditures of non-profit institutions serving households. We filled any gaps in the WDI with data from the IMF’s International Financial Statistics (IFS) or the country’s statistical offices. We complement the series in constant 2005 PPP USD with information from the WDI and the IFS in current and constant local currencies as well as current USD. We convert the current USD using market exchange rates. In the rest of the conversion we follow the same approach as with the micro data using consumption PPPs and CPI.

⁷⁰For the latest benchmark year, PovcalNet does not have a consumption survey for whole-China, so we used an income survey from the WYD.

We can re-arrange this to compute the Pareto coefficient a

$$a = \frac{1}{1 - \frac{\log(S_i/S_j)}{\log(H_i/H_j)}} \quad (4.2)$$

We use the top 10% and top 20% shares to compute the Pareto coefficient a for every country-year observation.⁷¹ Next we compute the top 1% and top 5% shares by using this estimate of a and solving Equation 4.1 for S_i . Then we can easily construct the new quantile groups: P99-P100 is simply the top 1%. P95-P99 is the top 5% share minus the top 1% share. Finally, P90-P95 is the top 10% share minus the top 5% share.

The validity of our results obviously rests on the parametric assumptions. The estimation is relatively flexible, since we estimate a different Pareto constant for each country-year observation. Our choice of functional form is relatively standard in the literature, where it is argued that top tails approximately follow a Pareto distribution (Atkinson et al., 2011). However, recent work on historical top income shares in Africa suggests that the distribution might depart substantially from a Pareto in some countries, and some years for a given country (Atkinson, 2014). We use the smallest available decile groups in our data, but other authors typically fit a Pareto distribution over smaller top fractiles (see footnote 71). When updating the database beyond 2008 in future work, we will use percentile groups where possible, so we can assess the implications of this assumption.

4.2.5 The panel dimension of our data

We are also interested in changes of a given country-decile-group over time. Hence, the panel dimension of our data is crucial. Table 4.2 counts the number of countries by the number of benchmark years for which a country appears in the data. For example, for 58 countries (out of a total of 162), mostly mature economies and Latin American countries, we have the complete panel. When we consider changes between

⁷¹Atkinson (2007b) uses the top 5% and top 10% shares. Hlasny and Verme (2013) use top 10% and top 1%. The World Top Incomes Database (WTID) (Alvaredo et al., 2013b) uses top 0.1% and top 1% in most cases.

two benchmark years, we do not require observations in the years in between, i.e. there could be gaps in the panel. For 63 countries, we can consider changes between 1988 and 2008. As a robustness check, we also consider the period 1993 to 2008, for which we have 90 countries and particularly improve the regional coverage of Sub-Saharan Africa and Russia/Central Asia/South-East Europe.

4.3 The cross-sectional distribution over time

4.3.1 Summary statistics

Since we are interested in analysing the global distribution of income, a first question to ask is how much of the world is represented by the surveys included in our database (Table 4.1). Because high-income countries are more likely to have a survey which can be included in our dataset, our coverage is higher when measured in terms of GDP (in 2005 PPP USD) than in terms of population. Our data represent 95% of world GDP on average and more than 90% in all benchmark years. On average (and in all years since 1993), our data also cover 90% of the world's population.

There are, however, substantial differences across regions. The coverage of Sub-Saharan Africa and Russia/Central Asia/SE Europe has improved markedly, in particular between 1988 and 1993.⁷² Our coverage of the Middle East and Northern Africa region, appears to have declined, particularly in the most recent benchmark year and more so in terms of GDP than population.⁷³ In the latter part of the analysis we present robustness checks for the period from 1993 to 2008, because 1988 has such a poor coverage of Sub-Saharan Africa and Russia/Central Asia/SE Europe.

The study most comparable to ours is Milanovic (2012) who does not use national accounts and uses the 2005 PPP exchange rates. Our samples differ in a number of ways. First, Milanovic allows for more detailed quantiles in particular after 1998, such

⁷²The first observation that we have for Russia is for benchmark year 1993.

⁷³This appears to be driven by the dropping out of Iran and Tunisia in 2008, which together represent 26% of the region's GDP and 23% of the region's population in 2008. Coverage in this region remains low because we miss big countries such as Saudi Arabia or the United Arab Emirates which account for 17% and 10% of regional GDP respectively.

as ventile groups, whereas our sample is restricted to deciles. *Ceteris paribus*, we would thus understate within-country and global inequality. Second, the data sources are different, although in most cases there is a substantial overlap between the PovcalNet surveys and the surveys used by Milanovic. Third, in our sample construction we have been more restrictive by (1) restricting countries to have either income or consumption surveys for all years, (2) surveys to be within two years of a benchmark year and (3) consecutive surveys to be at least three and no more than seven years apart. As a result, we lost potential benchmark years for 35 countries. This is the reason why the number of surveys in our database and our coverage of world GDP and population in benchmark year 1988 are lower than Milanovic (Table 4.1). In other benchmark years coverage is comparable and sometimes even higher in our sample. Furthermore, we obtain a larger fraction of surveys within plus or minus one year of the benchmark year.

4.3.2 The Gini index of the global distribution of income

Table 4.3 presents our main results on the inequality in the global distribution of income calculated across the unbalanced panel of country-decile-groups. Compared with within-country distributions, we find a very high level of inequality as measured by the Gini index: between 70.5% and 72.2%. Figure 4.1 shows the distribution of national Gini indices in 2008. The maximum Gini is 65.8% (Seychelles) which is less than the global Gini. Across the five benchmark years, the national Gini has an average of 39.3%, ranging from 19.4% (Czech Republic, 1988) to 74.3% (Namibia, 1993), which is the only country-year in excess of 70%. The global Gini index has remained virtually unchanged over time. Changes between benchmark years have been around 0.5%, with the exception of the period 2003 to 2008, when the Gini decreased by 1.89% or 1.36pp.

We could easily derive standard errors for the Gini index in order to account for sampling uncertainty (i.e. the fact that we have used a sample rather than the population). However, as Anand and Segal (2008) argue, these standard errors would not be appropriate, because they assume that there exists a single global household survey

with a clearly defined sampling uncertainty. In contrast, we have combined a large number of national household surveys each of which has its own sampling uncertainty. As a result, plausible standard errors should probably be substantially bigger, making the observed changes insignificant.

In addition, we agree with Anand and Segal (2008) that the non-sampling uncertainty arising from the following issues is likely to be more important. First, there obviously exists measurement error both in the household surveys and the price surveys used to estimate PPPs. Second, the mixing of consumption and income surveys and the use of benchmark years lead to uncertainties which are hard to quantify. The same applies to the changing global coverage over time. Third, our estimates of global inequality are biased downward, because we ignore inequality within country-year decile groups (although we keep this bias consistent across countries and over time). This bias is likely to be around 0.5pp (see discussion on p. 151). Fourth, the assumption of a common price level within the country (with the exception of China, India and Indonesia) leads to a bias in the same direction. The implications of regional price differences for within-country inequality are discussed in Chapter 1 of this thesis. On the other hand, our robustness check for missing top incomes, which allocates the entire gap between national accounts consumption and the household survey mean to the top decile group, is probably an upper bound.

In addition to the statistical significance, we might be interested in how substantial these changes are. Atkinson (2003) provides a useful metric in terms of the required tax increase to bring about an observed change in the Gini index. If we assume no government absorption and an approximately linear tax and transfer system, moving from a Gini index of 72.15% in 1988 to 70.50% in 2008, would require all country-deciles contributing 2.29% of income to a common pool which is then shared equally.⁷⁴

⁷⁴Atkinson (2003) shows that with an approximately linear tax and transfer system, the after-tax Gini can be written as

$$NetGini = \frac{(1 - t) * GrossGini}{\frac{mean\ net\ income}{mean\ gross\ income}}$$

where the denominator is $(1 - government\ absorption)$. Assuming no government absorption (as done by Deaton, 2010) and treating the 1988 estimate as the before-tax Gini and the 2008 value as the after-tax Gini, we have solved for the tax rate. Assuming a positive government absorption (e.g. part of development assistance is wasted in the process), would raise the required tax rate.

This is substantially greater than the UN target which urges all developed countries to contribute 0.7% of Gross National Product (United Nations, 2013).

As a robustness check, we computed the global Gini also for the balanced sample of countries (Table 4.B.2). Depending on the start date, there are four balanced samples. In all benchmark years and across the four balanced samples, the Gini index is always within 1pp of the results for the unbalanced sample. If we exclude the first benchmark year, the difference lies within 0.5pp. All results agree that the global Gini fell between 2003 and 2008, and between 1993 and 1998. The increase between 1998 and 2003 also seems quite robust, while there is some uncertainty over the direction of change between 1988 and 1993. Overall, our main results hold up to confining the analysis to the balanced panel.

As shown by Table 4.B.1, our estimates of the global Gini index are greater than many previous estimates in the literature.⁷⁵ The studies listed there differ fundamentally in their methodology, such as the use of national accounts aggregates, the type of PPP exchange rates and the interpolation for missing years.⁷⁶ The closest study to ours is Milanovic (2012) who also uses only surveys, and applies the 2005 PPP exchange rates.⁷⁷ Our estimate of the global Gini coefficient is bigger than Milanovic's (2012), although the gap is falling over time, from 4.35pp in 1988 to 1.76pp in 2003. This difference can probably be explained by the fact that Milanovic (2012) tends to have a lower share of consumption surveys compared with our data.⁷⁸ The direction of

⁷⁵Table 4.B.1 mostly draws on a similar table in Anand and Segal (2008). In addition, we report estimates from Milanovic (2002; 2005; 2012), Bourguignon (2012), Chotikapanich et al. (2012) and Warner et al. (2012). Note that Anand and Segal erroneously refer to Milanovic (2005) as GE(0), when in fact it is GE(1). For Milanovic (2002; 2005), we use the full sample, because this is closest to our approach. The decomposition in that paper is only available for the sample of countries that is common across all the years.

⁷⁶The fact that we and Milanovic (2012) obtain higher Gini coefficients than the rest of the literature might be explained by the fact that we use 2005 PPPs which follow the EKS rather than the GK method. The GK approach leads to a downward bias in global inequality, as explained above.

⁷⁷Milanovic (2012) takes account of rural-urban differences in price levels for China, India and Indonesia in the case of the Gini index, but not for the GE(1) measure, where he only reports figures allowing for these price differentials for China.

⁷⁸For example, in 1988 one third of our surveys are consumption surveys compared with 21% in Milanovic (Table 4.1). As long as there are positive savings, consumption would be less than income, so by having a higher share of consumption surveys we add more weight at the bottom of the distribution thus increasing inequality. Another difference to Milanovic (2012), which could go in the opposite direction, is that we understate within-country inequality by restricting the fractile groups to ten, whereas Milanovic allows for more detailed quantiles. However, in 1988 the average

change between benchmark years is similar with the exception of the period between 1988 and 1993.

4.3.3 Alternative measures of inequality

The Lorenz curves for 1988 and 2008 (Figure 4.2) intersect, i.e. different inequality measures will rank the two distributions differently. In this section we present results from the Generalised Entropy (GE) and Atkinson indices. The Gini index attaches a particular weight to inequality at different points along the income distribution. With the exception of the GE(2) measure, all inequality measures considered here find a decline in global inequality between 1988 and 2008 (Table 4.3).

The GE(0) (or Theil-L, or mean log deviation) index is particularly sensitive to differences in shares amongst low incomes, whereas the GE(2) index is sensitive to differentials at the top of the distribution (Cowell, 2009). The GE(1) (or Theil-T) index is an intermediate case. According to the top-sensitive GE(2) index, inequality increased between all benchmark years between 1988 and 2003. The GE(2) is sensitive to extreme values (Cowell and Flachaire, 2007); e.g. removing the maximum observation in 2008 reduces the GE(2) by 22%. Furthermore, the strong increase in the GE(2) is not robust to removing the largest observations: Removing the top (top five) observation reduces the increase to 3.10% (0.88%) instead of 15.92% in the full sample. After removing the top eleven observations, inequality falls by 1.96%. The bottom-sensitive GE(0) index shows falling inequality between 1988 and 1998, and a marginal, but probably insignificant, increase from 1998 to 2003. This appears to suggest that between 1988 and 2003, inequality amongst lower incomes was falling whereas it increased amongst higher incomes. Between 2003 and 2008, there has been a fall across the board.

We have computed the Atkinson (1970) index for three levels of inequality aversion ε . The higher ε , the stronger is the aversion to inequality in the distribution of incomes and the higher the weight attached to lower incomes. For $\varepsilon = 0$, society is indifferent

number of data points per survey is 10.79 in Milanovic which is similar to our ten groups. After 1998 Milanovic has substantially more groups per survey, such as 19.6 in 2002 and 2005. Hence within-country inequality is unlikely to explain the differences at least in the early years.

to the degree of income inequality. With $\varepsilon = \infty$, only the position of the poorest group matters. According to all three levels of ε , inequality is highest in 1988.⁷⁹ A(1) and A(0.5) agree on the relative rankings of the benchmark years (from lowest to highest inequality: 2008, 1998, 2003, 1993, 1988). For A(2), which is the highest level of inequality aversion considered here, inequality is higher in 2008 than 1998, and it does not change between 2003 and 2008 (at least to two decimal points). Furthermore, $\varepsilon > 2$ would show increasing inequality between 2003 and 2008, in contrast to all the other measures reported here.

All the measures considered so far are relative measures of inequality. An absolute measure of inequality is invariant to an equal absolute (or additive) change. A doubling of all incomes leaves the relative Gini unchanged, but it clearly changes absolute income differences, which is captured by an absolute measure. The absolute global Gini index is defined simply as the product of the global mean (panel H) and the relative global Gini index (panel A of Table 4.3).⁸⁰ In contrast to the relative Gini index, the absolute Gini index (panel F of Table 4.3) has risen substantially over this 20-year period, by some 20%.

4.3.4 Regional inequality and between-country decomposition

The Gini index calculated across all individuals living in a region is highest in Latin America and Sub-Saharan Africa. The mature economies have seen a strong increase in the last benchmark year. Inequality in China has risen strongly between 1988 and 2008, by more than 10pp.⁸¹ The increase in India has been much more moderate. The

⁷⁹The level of inequality aversion can be interpreted as follows (Atkinson, 1975). Consider two people who are identical except for one having twice the other's income. Consider a transfer which takes away \$1 from the rich and gives a proportion x to the poorer person. An inequality aversion of 2 implies that we would accept a transfer in which only \$0.25 reach the poorer person for every \$1 taken from the rich. For $\varepsilon = 1$, this corresponds to \$0.50 and for $\varepsilon = 0.5$, we would require \$0.71.

⁸⁰There are many formulations of the Gini index. According to one of these formulations (e.g. Anand, 1983 and Lerman and Yitzhaki, 1984), it is defined as the covariance between incomes and their ranks, divided by the mean (van Kerm, 2010): $Gini(X) = -2Cov\left[\frac{X}{\mu(X)}, (1 - F(X))\right]$, where $F(X)$ is the cumulative distribution function and $\mu(X)$ is the mean. The absolute Gini is defined as $AbsGini(X) = -2Cov[X, (1 - F(X))] = -2Cov\left[\frac{X}{\mu(X)}, (1 - F(X))\right] * \mu(X)$. We use Philippe van Kerm's `sgini` Stata routine to compute the absolute Gini index.

⁸¹It should be noted that our estimates of the regional Gini indices for China (and India) are based on 20 data points per year. Nevertheless, our estimates are not too dissimilar from other studies. Note

Gini index for Sub-Saharan Africa increased by approximately 5pp between 1993 and 2008. Within Middle East and Russia/Central Asia/SE Europe, inequality appears to have fallen. Inequalities within Latin America and Other Asia have remained virtually unchanged with some ups and downs in the intervening period.

We present a decomposition of the GE-class measures, which are, in contrast to the Atkinson and Gini indices, additively decomposable.⁸² We concentrate on the GE(0) index, because interpreting the within-group component as the residual inequality after equalising average incomes across countries, is only correct for this index out of the GE-class (Anand and Segal, 2008).⁸³ The between-country contribution has declined over this 20-year period, suggesting that countries, weighted by their populations, have become more similar.⁸⁴ In 2008, equalising mean incomes between countries while keeping the within-country distributions unchanged, would reduce global inequality by approximately 77%.⁸⁵ Alternatively, equalising all incomes within each country would reduce global inequality by 23% only. Compared with other studies in this literature, our between-country contribution is substantially greater (or, put differently, our within-country contribution is too small) (Table 4.B.1). This might be because we ignore within-decile inequality, whereas it is accounted for by at least some of the other

that we use consumption for China when we treat the rural and urban parts separately. Using NBS data and including adjustments for price differences, Ravallion and Chen (2007) report the following Gini indices for income (for all China): 1988: 29.50, 1993: 36.74, 1998: 35.37 and 2001: 39.45. Knight and Ding (2012) find a substantially higher level, but a similar trend, using CHIP income data: 1988: 39.8, 1995: 46.9, 2002: average of 46.2, 2007: 48.1.

⁸²The Atkinson index is decomposable by population subgroups, whereas for example the Gini index is not. However, the Atkinson index is not additively decomposable in the sense that it can be broken up into a weighted average of the within- and between-group inequalities (Bourguignon, 1979).

⁸³The GE(0) decomposition is in terms of population shares (see footnote 87), whereas the decomposition of GE(1) uses income shares. Redistributing income among countries in order to equalise average incomes, would change income (but not population) shares. A decomposition of the GE(2) index (not shown here) produces substantially smaller between-contributions compared with the GE(0) index. This suggests that at the top of the distribution, which is where the GE(2) measure is most sensitive, the between-country contribution is smaller, or in other words, countries are more similar to each other.

⁸⁴In an earlier draft, we also decomposed inequality by regions. With the exception of the GE(2) measure, the between-region contribution declined faster than the between-country contribution. This suggests that regions have become even more similar to each other compared with countries.

⁸⁵Treating China, India and Indonesia as single countries lowers the level of the between-country contribution. Furthermore, under this formulation, the within-contribution increases even more. This can be explained by the fact that inequality in whole-China increased faster than inequality within the rural and urban regions separately, precisely because average incomes between rural and urban China diverged. But differences between countries still explain 72% of global inequality using a GE(0) index, so the main conclusion remains unchanged.

studies, although this effect is likely to be small.⁸⁶ These differences are more likely to come from the change between 1993 and 2005 PPPs, which increased between-country inequality (Milanovic, 2012) (also see Section 4.1.3).

Inequality within countries has also increased in levels, not just as a share of global inequality (see panel D of Table 4.3). This is similar to Bourguignon (2011b), who also finds that while global and between-country inequality are declining, inequality within countries is increasing or at best unchanged. However, the between-country component, despite its relative decline, still remains by far the more important source of global inequality.

In order to further understand the driving forces of the changes in global inequality, we can also disaggregate the between- and within-components for the three most populous countries in the world; China, India and USA (see panel E of Table 4.3).⁸⁷ In 1988, the between-country components from these three countries alone accounted for almost 70% of global inequality (as measured by GE(0)). During the 20-year period, there has been a remarkable decline of the Chinese between-country contribution from 45% to 22%. This is explained both by the increase in average incomes in China and the decline in its share of the global population (from 26.6% to 21.7% of the sample). On the other hand, the Indian population share declined only very slightly (from 20.2% to 19.5%) and average incomes grew slower, and as a result the between-country contribution actually increased by 1pp. Furthermore, since 2003, the Indian between-component exceeds that of China, which follows from the much higher growth in average incomes in the latter (see panel H of Table 4.3).⁸⁸ The US between-contribution is negative

⁸⁶In a slightly different dataset of percentiles (see p. 151), moving from decile to percentile groups increases the within-contribution by 1.4pp.

⁸⁷The GE(0) measure can be decomposed between G population subgroups as

$$T_L = \sum_{g=1}^G \frac{N_g}{N} \cdot T_{L,g} + \sum_{g=1}^G \frac{N_g}{N} \cdot \ln \left(\frac{\bar{Y}}{\bar{Y}_g} \right)$$

where N is total population, $\bar{Y}$ is global mean, $T_{L,g}$ is GE(0) inequality within group g , and N_g and $\bar{Y}_g$ are subgroup population and mean respectively. The first sum is the within-component and the second refers to the between-component. The two sums can obviously be aggregated by different countries or regions. Note that the China and India contributions sum over the respective contributions from the rural and urban parts, which are treated as two separate countries.

⁸⁸The table shows the average income for all of China and India. By contrast, the country decomposition treats the rural and urban parts as separate countries, so it is the difference between (say)

because its average income exceeds the global average. The within-contributions increased for all three countries, particularly for China, as we already discussed for the regional Gini indices.

4.3.5 Growth incidence curves

Panel G of Table 4.3 shows the growth in incomes for a number of global percentile groups. The group which has grown fastest is the one between the 50th and 60th percentile (growth rate of 71.74% over 20 years), followed by the P40-P50 group (70.62%) and the global Top 1% (64.80%). Perhaps a more useful way to illustrate this pattern is through a variant of the global growth incidence curve (GIC).⁸⁹ It compares for instance the mean income of a given fractile group (e.g. the bottom 10% or the top 1%) in (say) 2008 with the mean income of the same fractile group in 1988. This is shown in Figures 4.3a, 4.3b, and 4.3c where the y-coordinate is simply the (annualised) growth rate between these two observations. A downward- (upward-) sloping GIC implies that economic growth has an equalising (disequalising) effect on the income distribution, i.e. it is pro-poor (pro-rich). These are anonymous GICs because the same fractile

the rural and global mean which matters. In 2003, the Chinese rural average exceeded its Indian equivalent for the first time. On the other hand, comparing urban averages, China exceeds India in all benchmark years. The rural incomes have a much stronger impact on overall inequality because they have a higher population share (in 2008, 12.4% for rural China and 13.8% for rural India). In both countries, the rural population share declined, while the urban share increased. In addition, rural average incomes grew slower than urban average incomes. In China, both urban and rural between-components declined although the rural component declined faster than the urban component. By contrast, in India, the rural contribution declined but the urban component increased. This is because in China, the growth in the urban average was fast enough to counteract the effect from the increasing population share, while this is not true of India.

⁸⁹The original GIC, as defined by Ravallion and Chen (2003), shows the growth rate in incomes for the same percentile (e.g. the 10th percentile of the global distribution) between the initial and final period. This can be obtained by inverting the cumulative distribution functions of the initial and final periods and taking the growth rate. Because it is derived from the cdf, this version of the GIC offers a straightforward test of first-order stochastic dominance. In contrast, we compare the mean income of the same fractile group (e.g. the bottom 10%) over time. The rationale is an important one when there is a distributional change within a given fractile group. To be precise: The original GIC evaluated at 100 percentiles fails to show the growth in incomes and the share if 100% of income gains are concentrated within the top 1%. Because the highest income group considered is the 99th percentile, the original GIC will show zero growth throughout. Of course, one could draw a GIC over (say) 1000 points instead. But this would make movements very erratic and broad trends harder to ascertain. Our averaging procedure could be seen as a smoothing technique. Not surprisingly, the two approaches differ most in the tails: Using the Ravallion and Chen approach, growth in the top (bottom) group is lower (higher) than in Figure 4.3a, with the effect being bigger in the top tail. However, the overall conclusions are robust and the supine *S* shape remains very clear.

group could be composed of entirely different people in the two years.

Figure 4.3a shows the global GIC for the period 1988 to 2008. As we already saw from Table 4.3, growth was highest in the P50-P60 range.⁹⁰ From around the 75th percentile, growth is lower than the growth in the global average. For the top 1% of the global distribution, growth reverts to being higher than the average. This gives the GIC a distinct supine *S* shape, with two peaks, around the median and at the very top, and a trough around the 80-85th percentile. The GIC is everywhere above zero (i.e. the average real income of every ventile group was greater in 2008 than in 1988), although this is not robust to using more detailed fractile groups in which case it drops below zero around the 80th percentile group. Figure 4.3b shows the global GIC for the separate five-year periods between benchmark years. The GIC for 2003-2008 lies almost uniformly above the other periods suggesting that growth has been highest over this period. During 1988-1993 incomes declined particularly for observations between the 70th and the 85th percentiles. The gains for the median and the top have been particularly strong in the last 2003-08 period whereas the losses for the groups around the 80th percentile have been exceptionally high in the first (1988-93) period.

Panel H of Table 4.3 shows the growth in average income for the different regions of the world. Not surprisingly, China is the region with the strongest growth, average incomes tripling between 1988 and 2008. In terms of 20-year growth it is followed by Russia/Central Asia/SE Europe and other Asia. In Russia/Central Asia/SE Europe, 15-year growth rates (i.e. 1993-2008) are also strong which might be a more appropriate time period given the low coverage of regional GDP and population in 1988. The only region with an income decline is Middle East and North Africa, although this might be affected by changes in sample coverage.⁹¹ Latin America has grown at a (marginally) lower rate than the global average. The mature economies have grown at a faster rate than India, although this is not true for all five-year periods. Sub-Saharan Africa almost did not grow at all between 1993 and 2008, which is the period when regional

⁹⁰The GIC shows annualised growth rates, whereas Table 4.3 shows total growth rates over this period.

⁹¹All calculations are done across the unbalanced sample of country-deciles.

survey coverage has been quite high for this region.

Figure 4.3c shows the 20-year GICs for the five regions for which regional coverage has been consistently high over this period.⁹² With the possible exception of the top 5% in Latin America, the GICs are everywhere above zero. Growth appears strongly pro-rich in China and less so in the mature economies and India, whereas the GIC is flat for Other Asia and displays no clear direction for Latin America.

While the global GIC showed relatively large gains for the portion of the distribution around the median, it is important to bear in mind that these gains were measured in relative (percentage) terms. But precisely because global income inequality is extremely high, and incomes at the top are several orders of magnitude greater than incomes at the median (in 1988, the average per capita income of the top 1% was close to \$PPP 39,000 while the median income was approximately \$PPP 600), the absolute gains are much greater for higher percentiles. Figure 4.3d (drawn on a logarithmic y-axis because the differences in absolute gains are so large) shows that the average per capita income for the top 1% increased by \$PPP 25,000 between 1988 and 2008, while the absolute gain at the global median was only \$PPP 400. The absolute gains among the poorer percentiles were even less. The overall outcome was thus that 44% of the increase in global income between 1988 and 2008 went to the top 5% of world population.⁹³

4.3.6 The regional composition of the global distribution

Figure 4.4 shows how the global distribution of income has changed over time. Income growth is shown by the rightward movement of the income distribution. The 1988 distribution appears to have two peaks, one around \$PPP 400 and another around \$PPP 10,000. In 2008, the second peak has disappeared and there appears to be more mass around the \$PPP 3,000 mark. As implied by the almost universally positive five-year GICs (Figure 4.3b), we can observe a rightward movement of the global distribution

⁹²The GICs by region are evaluated at decile groups with the top decile being split into two ventiles. This is because for China and India, which are regions by themselves, we have at most 20 observations.

⁹³The top 5% are defined relative to the population covered by the sample. These are, of course, not necessarily the same country-deciles (nor people) who were in the top 5% in 1988. We return to this issue in Section 4.5 where we discuss (quasi) non-anonymous GICs.

in every individual five-year period with the most striking development being the expansion of the proportion of the global population with incomes between \$PPP 750 and \$PPP 6,000 (that is, between approximately \$PPP 2 and \$PPP 16 per day). That population has expanded from 1.16bn or 23% of world population in 1988 to almost 2.7bn or 40% of world population 20 years later.⁹⁴

In order to disentangle these changes further, Figure 4.5 shows stacked kernel densities by regions.⁹⁵ Not surprisingly, the growth in China has had a profound effect on the global distribution. The change in the overall shape of the distribution appears to be driven by the upward movement of the upper decile groups in China. Both China and India have moved up along the income distribution while Sub-Saharan Africa seems caught at the bottom. The upward movement of China, because of its magnitude in terms of population and amount of growth, is particularly well illustrated in the stacked kernel densities. In 1988, the Chinese population was symmetrically distributed atop of the mode of the global distribution (exclusive of China). In other words, China and the rest of the world had about the same modal income. With each successive five-year interval, the Chinese distribution shifted more to the right (towards higher income levels) so much that by 2008 about four-fifths of the Chinese population has an income greater than the modal non-China global income.⁹⁶ The income mode in China is now clearly greater than in the rest of the world. It is this rightward movement of the Chinese distribution that has most contributed to the change from a two-peaked global distribution in 1988 to a single peak distribution 20 years later. This has largely happened because China has “filled up” the relatively hollow part of the global income

⁹⁴The total number of people refers to the entire world population, not only the population covered by our sample.

⁹⁵These charts have been created as overlaid (cumulative) kernel densities. Because the last density is shown on top we proceed in reverse order: The first density to be plotted is the global density including all regions (which is the same as Figure 4.4). Second, we plot the density for all regions, except China. We proceed by removing one region at a time until we end up with only the mature economies. The area under the global density is 1. The other residual densities are scaled down according to the regional population share in a particular year. For instance, in 2008 the second density we plot is scaled down to $0.7828 = 1 - x$, where x is China’s population share in 2008. We use an Epanechnikov kernel and the default bandwidth, selected optimally by Stata. The bandwidth is allowed to vary for different years, but for every benchmark year it is the same across the different cumulative kernel densities.

⁹⁶In 2008, 77.2% of Chinese had incomes exceeding \$PPP 650.

distribution between \$PPP 2,000 and \$PPP 6,000.

Figure 4.6 focusses on the change in the regional composition of the global income distribution between 1988 and 2008. The chart shows the regional composition of the population in each ventile of the global distribution.⁹⁷ As before, we can see a clear upward movement of China. The top decile group in China reaches as far as the 17th ventile (i.e. between 80th and 85th percentile) of the global distribution in 2008, whereas in 1988, the richest Chinese were only between the 65th and 70th percentile. Conversely, in 2008 China has entirely graduated from the bottom 5% of the world, while in 1988 it made up almost 40% of the population in that group.⁹⁸

As the bottom incomes in China have moved up the global distribution, Sub-Saharan Africa and to a smaller extent India have expanded their population shares in the bottom ventile. The distribution of Sub-Saharan Africa is very spread out with some decile groups (from South Africa and Seychelles) reaching the top 10% of the global distribution. India has not moved dramatically which is explained by the fact that its growth rate has been similar to the global average. The countries included in other Asia (not shown here) have also progressed out of the bottom ventile, but their movement into the upper parts of the global distribution has been less pronounced compared with China. This might be because growth was pro-rich in China, whereas it was more equally distributed in other Asia (Figure 4.3c).

The global 20-year GIC showed that fractile groups between the 75th and approximately 95th percentile grew slower than the global average (Figure 4.3a). In 1988, the percentiles between the 70th and the 85th (ventile groups 15, 16 and 17) originated primarily from the mature economies and Latin America, and to a smaller extent from the Middle East and North Africa (not all regions shown separately in Figure 4.6). By 2008, China and to a lesser extent Russia had moved into these percentiles, reducing the shares of the mature economies and Latin America, and the Middle East almost dropping out completely. It is these compositional changes which explain the

⁹⁷The ventile categories of course correspond to different absolute amounts in 1988 and 2008.

⁹⁸Of course, the caveat to this conclusion is that our dataset includes only 20 decile groups for China. If we had a continuous distribution, we would probably still observe some individuals in the global bottom 5%. However, the same caveat applies at the top end.

shape of the global GIC. The GIC does not track a particular fractile group, but rather compares the incomes of a given fractile in the *different* initial and final distributions. When comparing the top Chinese incomes in 2008 with the Latin American incomes in 1988, we obtain a below-average growth rate. This is despite the fact that the top Chinese incomes have grown substantially faster than the global average. This topic, the (quasi) non-anonymous GICs, which keep the composition of fractiles the same as in the original year, are discussed in Section 4.5 below.

4.4 Accounting for top incomes in the computation of global inequality

So far in the analysis we used our main cross-sectional sample and only the information contained in the surveys, and did not adjust survey means to national accounts. In this part of the analysis, we test the robustness of our conclusions to (1) anchoring to national accounts (i.e. imposing the same average income in national accounts and surveys), and (2) imputing for missing top incomes in surveys by distributing the excess income to the top decile group in each country using a Pareto imputation. Because this excess income seems too large in at least some country-years, we interpret our estimates of global inequality as an (extreme) upper bound.

Table 4.4 presents our results. Because we now use private consumption from national accounts we lose country-years for which we do not have national accounts information. This leaves us with 520 surveys across the five benchmark years instead of 565. In addition, we now treat China, India and Indonesia as single countries because we do not have separate national accounts information for rural and urban areas. And for China only we switch from consumption to income surveys (see p. 155). The baseline Gini for this new sample (row 1 in Table 4.4) is very similar to the Gini for the full sample (Table 4.3): the difference is less than 0.5pp except in 2008, when the new baseline Gini is 0.9pp lower than the one obtained from the full sample.

First, we scale the survey mean to equal national accounts consumption leaving

decile group shares unchanged (row 2 in Table 4.4). As discussed before, many papers anchor to GDP and we are the first to scale to national accounts consumption, which we argued to be more appropriate than GDP.⁹⁹ From row 3 onwards, we adjust country average incomes to the “censored” national accounts mean, which is defined as the maximum of national accounts private consumption per capita and the survey mean. This means that we replace survey means with private consumption only when the latter is greater, otherwise we keep survey means. In other words, national accounts consumption is censored below by the survey mean. We adopt our approach because we focus on the underestimation of top incomes in household surveys, thus assuming that surveys cannot overestimate, but only underestimate, household income. The global Gini is very similar whether or not we censor the national accounts mean (comparing rows 2 and 3 in Table 4.4), with the possible exception of 2008 (where the difference is almost 1pp).

Such “re-anchoring” obviously leaves within-country inequality unchanged. The entire effect on global inequality comes from the change in the between-component (and, in the case of the Gini index, indirectly from the change in the overlap component).¹⁰⁰ The between-component changes because the country-means change. A priori the direction in which we expect the global Gini to change is unclear.¹⁰¹ Comparing rows 2 and 3 with row 1 (Table 4.4), the Gini falls by approximately 1pp except in 2008 when it goes down by almost 2pp (for row 3). Intuitively, the reason for the downward change is that survey underestimation is greater in poorer (population-weighted) countries. In other words, poor countries appear less poor when we replace surveys with private consumption from national accounts.

⁹⁹Bourguignon (2012) uses GDP per capita and the 2005 PPP exchange rates. His estimates of the global Gini index are lower than the results in row 2, but follow a similar trajectory over time (see Table 4.B.1). Warner et al. (2012) additionally assume a parametric within-country distribution and obtain similar results to Bourguignon.

¹⁰⁰Milanovic and Yitzhaki (2002) show that the global Gini index can be decomposed by countries into a between-country term, the inequality within every country, its income share and an overlapping index of a particular country with the global distribution.

¹⁰¹Despite the many differences in these studies, we can probably conclude from Table 4.B.1 that estimates of the global Gini index are often lower when means are scaled to national accounts. In Milanovic (2005) [p. 118], scaling to GDP per capita *increases* the global Gini by 1.9pp in 1988 and by 0.2pp in 1993, but decreases it by 0.6pp in 1998.

Second, we assume that the distribution of the top 10% can be approximated by a Pareto distribution, which is similar to Atkinson (2007b).¹⁰² This “proportional adjustment with Pareto tail” allows within-country inequality to increase, while between-country inequality remains the same (since means are unchanged), and the overlap term remains the same or increases.¹⁰³ Consequently, we would expect the overall Gini, compared to the one from row 3, to be greater. This is indeed the case (see row 4) although the increase is quite moderate at less than 0.5pp (0.37pp on average). However, for other applications which focus on the top tail, such as the number of people above a certain income threshold (Atkinson, 2007b), the Pareto imputation is very important. Also note that a Pareto-based adjustment has a much greater effect on the global wealth distribution (Credit Suisse, 2013).¹⁰⁴ In the Appendix, we compare our Pareto coefficients to those in the World Top Incomes Database (WTID) (Alvaredo et al., 2013b), which offers more accurate evidence on top incomes (see Section 4.1.5). The proportional adjustment uses the income shares from the household surveys, because it scales up the mean while leaving (within-country) inequality unchanged. Given that household surveys underreport top incomes, it is not surprising that the Pareto coefficients from the proportional adjustment are too high compared to those obtained from the WTID. A higher Pareto coefficient implies lower inequality. Of course, we would not expect an exact correspondence because of differences between the two data sources (e.g. income definition or the unit of analysis).

Finally, we apply our “top-heavy adjustment with Pareto tail” (row 5 in Table 4.4),

¹⁰²Our approach differs from Atkinson’s in three respects: First, we use private consumption, whereas Atkinson uses GDP. Second, we keep survey means whenever they are greater than national accounts means (i.e. the censored national accounts mean), whereas Atkinson uses only national accounts information. Third, Atkinson assumes a Pareto distribution only for the top 5%.

¹⁰³The overlap component is likely to increase because the “elongation” of the distributions will tend to increase the “mixing up” of incomes of people from poorer and richer countries. The overlap component remains unchanged if the country-means are very far apart and the Pareto imputation only has a small effect. For example, if we have only Switzerland and Congo in the sample, the Pareto elongation of the top of the Congolese distribution may not yet place any Congolese fractiles within the range of the Swiss distribution (that is, all Swiss fractiles will have even higher incomes).

¹⁰⁴This particular adjustment uses a Pareto distribution in combination with information on top wealth holdings from the Forbes rich lists. The adjustment only stretches out the within-country distributions, leaving the country-means unchanged and hence is closely related to our “proportional adjustment with Pareto tail”. Their adjustment increases the global Gini from 87.6% to 90.5% (private correspondence with Anthony Shorrocks), which comes entirely from the within- and overlap-components.

where within-country inequalities are allowed to increase even further. This is done by allocating the entire (positive) gap between private consumption and the survey mean to the top decile group and applying (as in the previous case) a Pareto imputation. It should be intuitively clear that by increasing the share of the top decile group, we make income inequality greater and the Pareto constant lower. In the Appendix, we consider in more detail the size of the national accounts adjustment and also compare our Pareto coefficients and top shares to those in the WTID. Our national accounts adjustment seems excessive in some cases (e.g. in India, the survey mean is equal to only 53% of private consumption, see Table 4.B.3), an issue which we need to address more comprehensively in future work. Ignoring any definitional differences between national accounts and household surveys, the discrepancy between the two sources should probably also be allocated to lower incomes. Moreover, for some country-years the Pareto coefficients for the top-heavy adjustment are lower than what is observed in the WTID, i.e. the top income shares are too high. This is why we consider the “top-heavy adjustment with Pareto tail” (row 5) to be an extreme case. Compared with the “proportional adjustment with Pareto imputation” (row 4), the Gini increases by between 4.6pp and 7.8pp (6.17pp on average), with the most dramatic change occurring in 2008.

In order to test the sensitivity of the global Gini coefficient to these low Pareto constants, we limit the bounds within which Pareto constants can lie. We consider two ranges: First, the range observed in the WTID (see panel C of Table 4.B.4), i.e. $a = [1.29, 3.62]$. Second, we chose tighter (essentially arbitrary) bounds such that $a = [1.5, 3]$.¹⁰⁵ In both cases, the Pareto constants which are below the lower (above the upper) limit are changed to the lower (upper) limit. Using these revised values of a , we derive the size of the national accounts gap and in the final step compute the revised shares and average fractile incomes.¹⁰⁶ It seems that the main effect of restricting the Pareto constants is to increase some of the very low minimum values observed in

¹⁰⁵Based on tax data for the UK and US, the Pareto constant for income typically lies between 1.5 and 2.5, and is somewhat lower for wealth (Cowell, 2009, p. 93). Clark (1951) reports a large collection of Pareto coefficients which range from 1.13 (India in 1923) to 2.46 (New Zealand in 1937).

¹⁰⁶The national accounts adjustment ϵ is given by

Table 4.B.3, while the censoring from above has very little effect.¹⁰⁷ Table 4.4 (row 6) shows that limiting a to the range in the WTID, reduces the global Gini by between 0.22pp and 0.76pp compared with row 5. For the tighter limits (row 7), the differences lie between 0.78pp and 2.2pp. The three series track each other well.

In a second robustness check, we use the top 10% shares observed in the WTID to eliminate some of the extreme values. With the top-heavy adjustment, the (un-weighted) average income share of the top 10% is 46%. In India and Latin America, the average exceeds 50%, while in China it is 32%. In the WTID, the top 10% share ranges from 14% (Mauritius in 2005) to 46% (USA in 2010), with an average of 32% (approximately France in 1990s). We decided to censor the top 10% share at 50% and 40%. Hence we replace any excessive top 10% shares by the cut-off and re-compute the national accounts adjustment.¹⁰⁸ In other words, this censoring removes some of the very high levels of inequality generated by our imputation. This adjustment only eliminates top shares which are too high (i.e. a type of one-way censoring), which is where we saw the biggest effects when limiting the Pareto constant. In Table 4.4 (row 9), censoring the top 10% share at 40% gives the midpoint between the baseline and the top-heavy Pareto, while censoring at 50% (row 8) gives values closer to the top-heavy results. Furthermore, when using the more restrictive bounds of 40%, the decline between 1988 and 2008 is quite similar to the baseline with only survey data.

In the two robustness checks we first adjusted the Pareto constant or the top 10% share and second computed the national accounts gap. It might appear more straightforward to adjust the national accounts gap directly. The justification for our approach

$$\epsilon = \frac{1}{10} \left[\frac{\bar{Y}_9}{2^{\frac{a-1}{a}} - 1} - \bar{Y}_{10} \right]$$

where $\bar{Y}_9$ and $\bar{Y}_{10}$ are the average incomes of the 9th and 10th decile group observed in the household survey.

¹⁰⁷When we do not impose an upper bound on the Pareto constant (i.e. a lower bound on inequality), the global Gini is within 0.1pp of the value for the two-way censoring.

¹⁰⁸The national accounts adjustment ϵ is given by

$$\epsilon = \left[\sum_{i=1}^9 \frac{\bar{Y}_i}{10 \times (1 - s_{10}^*)} \right] - \bar{Y}$$

where $\bar{Y}_i$ is average income of decile group i , $\bar{Y}$ is average income in the survey, and s_{10}^* is the censored top 10% income share.

is that we had some guidance on what might be a sensible range of values from the WTID. Note that in both robustness checks, for the observations which are censored, the country-mean no longer equals the national accounts consumption. Hence these checks affect both within- and between-country inequality.

Figure 4.7 summarises the global Gini index for the main specifications. We show the relative and absolute baseline Gini indices from Table 4.3 and rows 2, 5, 6, 7, 8 and 9 from Table 4.4. The global Gini calculated from survey data alone (Table 4.3) is reduced by up to 2pp (1.1pp on average) when we replace survey means by private consumption from national accounts (thus allocating the gap proportionally across national income distributions). Allocating the positive national accounts gap to the top decile group and applying a Pareto imputation, increases the global Gini by between 4.18pp and 6.22pp compared to the baseline. The truth probably lies somewhere in between. This gap is reduced to between 3.88pp and 5.65pp when limiting the Pareto constants to the WTID range. For the most conservative range of Pareto constants, the top income adjustment raises the global Gini by between 3.31pp and 4.65pp, which is still a substantial upward revision of the global Gini. Restricting the top 10% share to 50% produces results very close to more conservative Pareto range. Using a more restrictive 40% cut-off brings the global Gini to the mid-point between the top-heavy Pareto and the baseline. These last two results suggest that a substantial part of the global distribution is affected by moving from a 50% to a 40% income share for the top 10%. From the existing evidence, a top 10% share in excess of 40% is high but not unprecedented. In sum, the global Gini we estimate ranges from 68.6% to 75.9% in 2008. The sensitivity checks (with the exception of the 40% censoring) suggest that the truth is closer to the upper bound, although this is of course no definite proof.

A further result of this exercise is that most of the top-heavy adjustments, with the exception of the 40% censoring, reduce the change in the global Gini index between 1988 and 2008, compared with Table 4.3. For the most extreme adjustment, the global Gini changes by only 0.6% over these 20 years. This is because the size of the national accounts adjustment increased from an average of 19% in 1988 to 25% in 2008 (see

Table 4.B.3). When we allocate this rising gap entirely to the top tail, we obtain an increasing within-country, and ultimately global, inequality. Limiting the Pareto coefficient, or the top 10% share, and thus the national accounts gap, increases the observed change. Before, we argued that the change in the global Gini index observed for the full sample (and using incomes directly reported in the surveys), was probably not robust to plausible standard errors. This robustness check further supports a more cautious view about the decline in global inequality: if indeed surveys tend to underreport incomes at the very top, it could well be that global inequality, measured by the Gini index, has not gone down during the 20-year period considered here.

It can be shown that the reasons why the decline in global inequality disappears are a strong increase in the inequality amongst the top group and their share of total income. Using the Alvaredo (2011) decomposition, we can decompose the top-heavy global Gini index into the Gini indices in the bottom and top groups, weighted by their income and population shares.¹⁰⁹ For the bottom 99% of the global distribution, the Gini coefficient declines by 3.3pp, similarly to the baseline. However, this decline is counteracted by a 5.6pp increase in the Gini amongst the global top 1%, and a 6.3pp increase in their share of total income. Therefore, the overall global Gini with a top-heavy Pareto declines by only 0.5pp.

Using the income reported in surveys we concluded above that the global income distribution had moved from a twin peak to a single peak. This also holds for the global distribution of income which adjusts for missing top incomes (Figure 4.8). Not surprisingly, the top-heavy Pareto imputation stretches out the top tail and makes it thicker from around \$PPP 40,000. Further down the distribution, mass appears to shift from around \$PPP 3,000 to \$PPP 5,000, which are the top fractiles in the poor countries.

¹⁰⁹The decomposition is given by $G = G^{**}PS + G^*(1 - P)(1 - S) + S - P$, where G^* is the Gini coefficient for the non-top group, G^{**} is the Gini coefficient for the top group, and P and S are the population and income shares of the top group.

4.5 Changes over time: Who has grown fastest and slowest?

In this section we return to the issue of which country-decile-groups have contributed to the overall changes in the global distribution. The evidence on the changing regional composition (Figures 4.5 and 4.6) was a first step into this direction, although it did not say anything about the movement of individual country-deciles. This section tries to identify the particular country-deciles with the highest and lowest growth rates using a quasi-non-anonymous growth incidence curve.

The sample used in the main text includes all countries which are observed in 1988 and 2008, even if there are gaps in the intervening benchmark years. In the Appendix, we have replicated the results using all countries observed in 1993 and 2008, which improves the coverage particularly in Sub-Saharan Africa and Russia/Central Asia/South-East Europe (Table 4.2).

4.5.1 Quasi-non-anonymous growth incidence curves

In order to identify the winners and losers, we consider “income mobility profiles” (van Kerm, 2009), or “non-anonymous growth incidence curves” (Grimm, 2007; Bourguignon, 2011a). When applied to individual-record data, the distinction into anonymous and non-anonymous GICs is straightforward: The standard anonymous GIC compares the incomes of say the 20th percentile in the initial and final period distributions. As long as there is some mobility in the distribution, the individuals at this percentile might be very different. By contrast, the non-anonymous growth incidence curve is the (non-parametric) regressions of (annualised) income growth against the rank in the initial distribution. Because this growth rate is obtained for each individual, it is non-anonymous, taking into account the joint distribution of initial and final incomes. However, in our case the unit of analysis are income decile groups of a particular country, so while we preserve the identity of a particular country-decile, these decile groups are defined over different people. Hence we refer to our charts as

“quasi-non-anonymous” growth incidence curves.

Figure 4.9 shows the quasi-non-anonymous GICs for 1988-2008 and 1993-2008. It plots the annualised growth over the next 20 (15) years against the normalised fractional rank in the 1988 (1993) *cross-sectional* global income distribution.¹¹⁰ In order to exploit both the cross-sectional and the panel dimensions of our data, each observation included in the quasi-non-anonymous GIC is ranked in the complete cross-sectional 1988 distribution (population-weighted) (not only amongst the 63 countries observed in 1988 and 2008). The main result is the curve labelled local polynomial, which is the fitted line from a kernel-weighted local polynomial regression. For comparison, we added a linear and a quadratic fit. For each time period, we present charts with and without the scatterplot.¹¹¹ The scatterplots show the wide dispersion of growth rates around the fitted line. Judging from the scale of the vertical axis, the dispersion seems to have increased, but this might be driven by one outlier close to the bottom ranks in 1993-2008. The fitted lines are shown on a more detailed vertical axis in the bottom panel for the 1988-2008 and 1993-2008 periods.

It is immediately apparent that the shape of the quasi-non-anonymous curves is very similar to the shape of the anonymous GIC (Figure 4.3a): they all display a supine *S* shape driven essentially by very slow growth around the 80th and 90th percentile of the global income distribution and local maxima around the median of the income distribution and for the very top. However, if we compare the 1988-2008 results, it is clear that the gains among the country-deciles that were in the top 1% in 1988 were less than the gains we obtain by simply comparing income levels of the top 1% in 2008 and 1988. This is not surprising because not every country-decile that was in the top 1% in 1988 managed to have high growth over the next 20 years. Similarly,

¹¹⁰We also repeated the analysis using an alternative mobility measure on the vertical axis such as the change in ranks, which produced a similarly-shaped curve. Fractional ranks are derived from a smooth cumulative distribution estimator which ensures that the mean rank is 0.5, estimated using the `fracrank` Stata routine by Philippe van Kerm.

¹¹¹There are two reasons why observations are not equally distributed along the horizontal axis, and appear concentrated amongst the upper ranks. First, the ranks are computed in the cross-sectional distribution, whereas the scatterplot includes only those country-deciles which are observed in 1988 and 2008, many of which are from rich countries. Second, country-deciles are weighted by population size in calculating the ranks, whereas this information does not show up in the scatter points.

some country-deciles that were not among the top 1% in 1988 and have exhibited high growth are now (in 2008) in the top 1%. We find the equivalent result for the 1988 poorest country-deciles whose growth rates were higher than what we found from the anonymous GIC. Furthermore, some of these differences might arise from restricting the sample to those countries which are present both in 1988 and 2008. Over the 1993-2008 period, growth was highest around the 60th percentile. The shape of the fitted line also seems to have changed with some high growth rates amongst the very bottom 1993 ranks.

Over the 1988-2008 period, growth was highest for those country-deciles around the 40th percentile of the 1988 global distribution, and lowest around the 85th percentile. However, these differences are small (approximately 2pp of annualised 20-year growth rates) and a reasonable standard error would probably reduce these differences to insignificance. Groups that were most successful come overwhelmingly from China and India; groups that were least successful are predominantly from the mature economies. Three quarters of the population that was between the 36th and 45th global percentile in 1988 belonged to country-deciles from China and India, generally around the middle of national income distributions. If we also include other Asia, the percentage of people belonging to this most successful group reaches 90%.

In contrast, the country-deciles between the 81th and 90th percentile in 1988 are overwhelmingly from mature economies, and come from the lower halves of their national income distributions. Out of 420 million people belonging to this group, 364 million are from the mature economies (or differently, 135 out of 165 country-deciles). Even when we exclude from the mature economies those that in 1988 were Communist, the share of the “traditional” rich Western economies among this group is still very large at 78% of people.

4.5.2 Most successful and least successful

4.5.2.1 Ranked by growth rate

We next look directly at the 20 biggest gainers and losers (country-decile-groups) over this 20-year period (Table 4.5).¹¹² In the first two columns, we rank country-deciles according to the average of the (annualised) 5-year growth rates between 1988 and 2008.¹¹³ In columns 3 and 4, we rank according to (annualised) growth in decile group income between 1988 and 2008, which is the same as the vertical axis in Figure 4.9.¹¹⁴ Incidentally, the two rankings do not make much difference to the 20 country-deciles which have grown the most.

Between 1988 and 2008, all but one of the top 20 country-deciles have grown in excess of 5% annually, which means that their incomes have increased by at least 2.65 times. For the decile groups that have grown at around 8% annually, real incomes have been multiplied by a factor greater than 4.5. Almost all urban Chinese decile groups are amongst the 20 fastest growing decile groups in the world.¹¹⁵ Moreover, they are neatly ranked with the highest decile groups having grown fastest. This illustrates the already observed disequalising pattern of growth in China. Rural China has grown slower but the top two rural decile groups are also among the 20 winners, and again are ranked, with the rural top decile group growing faster than the second highest rural decile group. In addition, some decile groups in Latin America and some mature economies such as Ireland and the UK, are amongst the top 20. With the exception of some Latin American decile groups, the bottom 20 are almost exclusively from Eastern Europe.¹¹⁶ Estonia and Lithuania have a low average growth rate (columns 1 and 2), but comparing 1988 and 2008 incomes they are not in the bottom 20 (columns 3 and

¹¹²Decile groups are numbered 1 to 10, with 1 being the bottom group.

¹¹³This excludes any countries which have no 5-year intervals. Cyprus (for 1988-2008) and Niger (for 1993-2008) are observed in the initial and final years, but have no 5-year growth intervals in between.

¹¹⁴The bottom decile group in Japan declined by an annualised rate of 3.05%, which seems very high. Over the same period, average incomes in Japan (in our data) declined by 0.33% per year, whereas GDP per capita (at 2005 \$PPP) grew by 1.34% per year. Only in 2008 we were able to obtain micro data from LIS for Japan. Before that the data are based on tabulations, which might explain the substantial decline between 1988 and 2008 for the bottom decile group.

¹¹⁵This has nothing to do with population size because each country-decile is treated equally here.

¹¹⁶Since all these countries are EU members, they are included in the mature economies region.

4). This is because their incomes declined in the intervening period, e.g. between 1988 and 1993, but recovered by 2008.¹¹⁷ This is confirmed by Table 4.B.6, which ranks Lithuania in the top 20 according to growth between 1993 and 2008.

These results are quite different for 1993 to 2008 (Table 4.B.6). The high growth rates of the Chinese decile groups are now exceeded by some Sub-Saharan African, Eastern European and Russian/Central Asian countries. Of course, increased sample coverage could have played a role, particularly in the case of Sub-Saharan Africa, where countries might have previously been missing. The country-deciles that have lost most between 1993 and 2008 are mostly from East Africa and Latin America, with Eastern European decile groups having recovered. These results are consistent with the disastrous income declines in Eastern Europe between 1988 (the beginning of the reform process) and mid-1990s followed by, in some countries sharp, and in other countries, rather slow recovery.

4.5.2.2 Decomposition of growth rates into growth and redistribution

For both time periods, we have also decomposed the decile group growth rate into a growth effect (arising from the growth in the country mean), a redistribution effect (arising from a change in the decile share) and a residual component. This is motivated by the decomposition introduced by Datt and Ravallion (1992) in the context of the poverty headcount index. The share of total income received by a particular quantile (e.g. decile) group i in country c at time t can be written as $S_{i,c,t} = \frac{Y_{i,c,t}}{Y_{c,t}} = \frac{1}{Q} \frac{\bar{Y}_{i,c,t}}{\bar{Y}_{c,t}}$, where $Y_{c,t}$ is total income and $Y_{i,c,t}$ denotes total quantile income; $\bar{Y}_{c,t}$ and $\bar{Y}_{i,c,t}$ are the respective means; and Q is the number of quantile groups (e.g. 10 in the case of decile groups). It can be shown that the (20-year) growth rate in the quantile mean can be decomposed into the growth rate in the country mean, the growth rate in the quantile share and a residual term which is the product of the two¹¹⁸

¹¹⁷Given the economic upheavals during the break-up of the Soviet Union, measurement error over the period 1988-1993 might have been more pronounced than during other periods.

¹¹⁸The decomposition can be obtained as follows: First, re-write the numerator and denominator using $\bar{Y}_{i,c,t} = QS_{i,c,t}\bar{Y}_{c,t}$, and equivalently for $\bar{Y}_{i,c,t+1}$. Second, add and subtract $\frac{S_{i,c,t}\bar{Y}_{c,t+1}}{S_{i,c,t}\bar{Y}_{c,t}}$, and $\frac{S_{i,c,t+1}\bar{Y}_{c,t}}{S_{i,c,t}\bar{Y}_{c,t}}$, and collect terms.

$$\frac{\bar{Y}_{i,c,t+1} - \bar{Y}_{i,c,t}}{\bar{Y}_{i,c,t}} = \frac{\bar{Y}_{c,t+1} - \bar{Y}_{c,t}}{\bar{Y}_{c,t}} + \frac{S_{i,c,t+1} - S_{i,c,t}}{S_{i,c,t}} + \frac{\bar{Y}_{c,t+1} - \bar{Y}_{c,t}}{\bar{Y}_{c,t}} \times \frac{S_{i,c,t+1} - S_{i,c,t}}{S_{i,c,t}}$$

In columns 5, 6 and 7 of Table 4.5 (and Table 4.B.6), we show the results of this decomposition, where we divided by the decile group growth rate, so the contributions add up to 100%.¹¹⁹ Because the residual term has no particular interpretation, the usefulness of this decomposition obviously depends on the relative size of this term.¹²⁰ Many of the residuals in Table 4.5 (and Table 4.B.6) are quite large and reduce the usefulness of this decomposition. Nevertheless we can draw some conclusions.

Between 1988 and 2008 (Table 4.5), the gains by the urban Chinese decile groups appear to be mostly driven by growth in average income, although there was some redistribution towards the top decile group. For the Latin American decile groups (with the exception of Guatemala) at the top and bottom of the distribution of growth rates, redistribution seems more important. For some countries in this region, redistribution was towards the top (e.g. Mexico) or away from the bottom (e.g. Honduras, Paraguay, Bolivia) while in others it was towards the bottom (e.g. El Salvador, Costa Rica).¹²¹ The Eastern European decile groups amongst the bottom 20 are explained primarily by negative or slow economic growth, although in Latvia anti-poor redistribution seems to have played a role.

¹¹⁹A positive contribution means that the source (growth/share/residual) works in the direction of the decile group growth rate, which might be positive or negative. On the other hand, a negative contribution implies that the source works in the opposite direction. We show the annualised decile group growth rates, but the decomposition is based on the 20-year growth rates.

¹²⁰The residual term is the result of an index number problem which arises because one period is arbitrarily chosen as the reference period. We could derive two decompositions, which anchor to the initial and final period respectively. Kolenikov and Shorrocks (2005) show that the residual term drops out when they take the average over these two decompositions, and that this is based on a Shapley-decomposition. In our case, this would work with absolute changes in the decile group mean, but it does not work with the growth rate. Datt and Ravallion (1992) consider year-to-year changes in the poverty headcount, whereas we consider changes over the longer term, which might explain why the residuals are smaller in their application.

¹²¹For Bolivia-1 and Honduras-1, the share contributions have positive signs which implies that they work in the same direction as the decile group growth rates, which are negative. On the other hand, the growth contributions have negative signs, i.e. the growth in the mean had an opposing (in this case positive) effect on the decile group growth rate. For Paraguay-1, both contributions work in the direction of the decile group growth rate (i.e. the country-mean also declined), but the share contribution is greater than the growth contribution.

In the 1993-2008 period (Table 4.B.6), economic growth explains the strong performance of Lithuania at the top, while the other decile groups at the top tell a more mixed story. At the bottom, Tanzania's decile groups have lost out because of a decline in country average incomes. For the 5th to 10th decile groups in Kenya, both growth and redistribution seem important, whereas in Bolivia it is primarily due to redistribution away from the poor.

4.5.2.3 Positional gains in the global distribution

Another simple way to evaluate the success of various country-decile-groups relative to the rest of the world is to compare their positions in the global income distribution. This can obviously be done for every country-decile and every year. In Figure 4.10, we do it for several selected countries (2008 values are always drawn as a solid line, 1988 values as a dashed line). The top left panel illustrates the already discussed remarkable upward mobility (in the global income distribution) of Chinese rural and urban decile groups. In 2008, the Chinese top urban decile group is at the 83rd global percentile while 20 years earlier it was at the 68th. In other words, if somebody stayed in that same decile group in China over 20 years and his income followed the average growth path of the decile group, he would have leap-frogged, in terms of income, more than 900 million people worldwide. It is interesting that in 2008, the Chinese top *rural* percentile is at a higher global position than the Chinese top *urban* percentile was in 1988. The same development, although less dramatic, is illustrated by the improvements in the position of Indonesian rural and urban decile groups shown in the upper right panel of Figure 4.10.

But when we turn to the bottom left panel the situation is exactly the reverse. There we see that Nigeria's and Côte d'Ivoire's decile groups have almost uniformly slid down the global income distribution. It is only the top Ivoirian decile group that has managed to preserve its 1988 position. The bottom Nigerian decile group remains unchanged in the lowest global percentile. Finally, in the bottom right panel, we show Germany, Brazil and Ghana. The position of the German decile groups has remained

very high and displays little change. Brazil is an example of a reasonably fast growth across the income spectrum, and improvements in the distribution such that all its decile groups, except the highest, are now placed globally higher than they were in 1988. Ghana had a similar evolution between 1988 and 2003 (there are no data for 2008), except that it is the poorest decile group which has not seen any improvement in its position in the global distribution. This stands in sharp contrast to neighbouring Côte d'Ivoire: As a result of the growth in Ghana and the decline in Côte d'Ivoire (bottom left panel), Ghana in 2003 is richer than Côte d'Ivoire in 2008, thus reversing their rankings in 1993.

4.6 Conclusion

The chapter has provided evidence on the evolution of the global income distribution during a crucial period of accelerated globalisation and marketisation spanning the period from the end of Communist regimes in Eastern Europe and the Soviet Union to the beginning of the global financial crisis. In many respects, this might have been the most globalised period ever in history. It is very important to study what were its effects on the level and distribution of income among the world population. Our results confirm earlier findings that the level of global inequality remains high, with a Gini of around 0.7, and while inequality appears to have declined in the most recent years, these changes are probably not robust to plausible standard errors.

A robustness check whereby missing incomes, defined as the positive gap between per capita household consumption from national accounts and mean per capita income from household surveys, are allocated to the top 10% of individuals, and top incomes are supposed to follow a Pareto distribution, returns estimates of the global Gini coefficient which are about 5pp higher. Most of the increase is due to the allocation of the entire gap to the top decile group, not to the Pareto imputation of the top tail. A value of 0.75 may be considered an upper bound on the “true” global Gini. In effect, we assume that the two problems that have been discussed separately - namely the large, and in some cases, growing gap between national accounts consumption and means from household

surveys, and the realisation, due to the results obtained from fiscal data, that surveys tend to underestimate top incomes - are closely connected issues. We want to stress that this robustness check is very preliminary and that in future research we should attempt more sophisticated adjustments for missing incomes at the top.¹²²

The shape of the global income distribution has also changed during the 20 years considered here. In 1988, the global income distribution displayed a twin-peak shape which has since disappeared. This was mostly due to the high growth of China whose decile groups have “filled up” the area between \$PPP 2,000 and \$PPP 6,000 that was relatively “hollow” in 1988. The period has also witnessed a remarkable increase in what may be called a “global median class”, with incomes ranging from \$PPP 2 per capita per day to \$PPP 16 per capita per day: the share of the global population belonging to that group has increased from some 23% to 40%.

Particularly important is the shape of the anonymous and quasi-non-anonymous growth incidence curves for the period 1988-2008. They both show large relative gains around the median of the global income distribution that accrued mostly to the middle or upper-middle decile groups from Asia, and especially from China. By contrast, the lowest real income gains were registered around the 80-85th global percentiles where low income decile groups from the mature economies were overrepresented. A striking fact is that among the percentiles in 1988 that turned out 20 years later to have been the most successful part of the global income distribution, 90% of people came from Asia. Among the 1988 percentiles that 20 years later turned out to have been the least successful part of the global income distribution, 87% came from mature economies.

The chapter has offered another contribution. It constructed a new database, which (a) consists almost entirely of data derived from household survey micro data; (b) has

¹²²The national accounts gap appears too large in some cases. We have already tried to limit the extent of this adjustment but this was essentially arbitrary. We plan to allocate the survey-national accounts discrepancy to a bigger segment of the income distribution, such as the top 20%. The most rigorous approach would be to combine our data with evidence from the WTID. But this cannot be done yet for estimating global inequality because of the lack of widespread data availability. Evidence from global “rich lists”, such as the Forbes list provides information about the wealth of the extremely rich. In a developing country context, we may also be particularly concerned with incomes not reported to tax authorities. Data on illicit financial flows and wealth held in tax havens might provide some information (Zucman, 2013).

a consistent allocation of surveys to benchmark years; (c) keeps the household survey concept (income or consumption) constant over time for a given country; (d) uses only the most recent and most robust PPP exchange rates; and (e) allows a balanced panel analysis across the country-decile-groups. The latter aspect is particularly important since it has not received much attention in the literature. In other words, while we knew that China's growth rates were high and that the process was pro-rich, we could not directly compare the growth of China's top decile group with that from the US, Germany or Ghana. Indeed we find that all Chinese decile groups, both rural and urban, have improved their relative position in the global income distribution, at times dramatically by jumping by more than ten percentiles. The top urban Chinese decile group thus went from being at the 68th percentile in the world to the 83rd. By contrast, in some African countries the positions of all or almost all of their decile groups have gone down globally. China accounts for half of the top 20 most successful country-decile-groups in the period 1988-2008. Among the less successful country-deciles, the share of Eastern European and more recently African decile groups is preponderant.

In future research, it would be interesting to address the sensitivity of our results to the various assumptions we (and others) have made in deriving these estimates. In particular, one might consider the effects of (1) using a particular set of PPP exchange rates; (2) allowing for inequality within decile groups; and (3) the mixing of consumption and income surveys. Although our PPPs are substantially improved compared with earlier studies, they are still a source of uncertainty. A new set of PPPs will serve as a robustness check. To assess the effect of underestimating within-country inequality, we could fit parametric Lorenz curves to the decile group shares and estimate continuous within-country distributions, as suggested by Anand and Segal (2008) (using for example the Shorrocks and Wan (2008) procedure). From the estimated parameters we can then create a dataset for each country-year. In other words, we would create a global dataset where each observation represents the same number of persons, as would be the case in a truly global household survey.¹²³ The validity of our approach obvi-

¹²³The "sampling probability" (i.e. how many persons each observation represents) would have to be chosen so as to guarantee a minimum number of observations per country. Otherwise the smallest

ously depends on the parametric assumptions¹²⁴, and it could be argued that trying to recreate observations from grouped data is unnecessarily convoluted.¹²⁵ We did not consider this adjustment here, because we tried to stay as close as possible to the data and thus avoided imposing particular parametric assumptions. Furthermore, having a database of country-year decile groups eased the discussion particularly in the dynamic part. In a related dataset moving from decile to percentile groups changes the global Gini by only 0.5pp, so the effect of this underestimation is probably quite small.

Trying to assess the sensitivity of our results to mixing income and consumption surveys is more difficult. One possibility is to make an essentially arbitrary assumption about the differences between the two types of surveys and adjust them accordingly. For example, some authors have suggested scaling income surveys down by the amount of savings in national accounts (Deaton, 2001). Alternatively, one could estimate a relationship between income and consumption inequality for countries which have both types of surveys, and then extrapolate to the rest of the sample. Another possibility, which we prefer, is to estimate separate global distributions for income and consumption. Given that we selected the type of survey which maximised the number of benchmark years for each country such an approach would obviously reduce the global coverage of our database. For this reason, we did not adopt this approach here.

country might only have one observation. Alternatively, we could of course stratify the sample by country-year and weight each observation accordingly.

¹²⁴There exists a substantial literature which fits parametric income distributions to fractile shares in order to account for within-fractile inequality (see footnote 47 for more details). Recently, Chotikapich et al. (2012) use a three-parameter beta distribution. Pinkovskiy (2013) shows that allowing for an arbitrary within-country distribution (given fractile shares) makes the direction of change in the global Atkinson index ambiguous.

¹²⁵For example, we have converted ventiles and percentiles into decile groups. We have more detailed fractile groups only in recent years. PovcalNet fits a Lorenz curve to the underlying more disaggregated data and reports decile groups. Going back to the original micro data is not possible, since most information is only released in grouped form.

4.7 Main Tables and Figures

Table 4.1: Sample summary statistics

	Benchmark year					Total
	1988	1993	1998	2003	2008	
Number of surveys	75	115	121	133	121	565
<i>Milanovic (2012)</i>	<i>103</i>	<i>122</i>	<i>124</i>	<i>122</i>		
Years between survey year and benchmark year (% by benchmark year)						
-2	12.00	9.57	9.09	7.52	7.44	8.85
-1	26.67	18.26	14.88	18.80	11.57	17.35
0	29.33	34.78	41.32	30.08	65.29	40.88
1	18.67	20.87	18.18	21.05	11.57	18.05
2	13.33	16.52	16.53	22.56	4.13	14.87
Within +/- 1 of benchmark	74.67	73.91	74.38	69.93	88.43	
<i>Milanovic (2012)</i>	<i>71</i>	<i>67</i>	<i>70</i>	<i>66</i>		
Income vs. Consumption surveys (% by benchmark year)						
Consumption	33.33	46.09	48.76	57.14	55.37	49.56
Income	66.67	53.91	51.24	42.86	44.63	50.44
<i>Milanovic (2012): % consumption</i>	<i>21.36</i>	<i>42.62</i>	<i>50.81</i>	<i>54.10</i>		
GDP (2005 PPP USD) (% of regional GDP represented in the database)						
World	90.60	96.97	96.51	95.93	92.97	94.60
<i>Milanovic (2012): World</i>	<i>96</i>	<i>95</i>	<i>96</i>	<i>98</i>		
Mature economies	95.65	99.90	99.83	98.36	96.87	98.12
China	100.00	100.00	100.00	100.00	100.00	100.00
India	100.00	100.00	100.00	100.00	100.00	100.00
Other Asia	90.51	96.12	97.99	97.31	97.05	95.80
M. East & N. Africa	52.13	55.40	45.39	48.48	22.27	44.73
Sub-Saharan Africa	21.94	82.90	78.31	81.83	77.37	68.47
L. America & Caribbean	94.77	98.27	99.02	98.85	98.35	97.85
Russia, C. Asia, SE Europe	50.68	94.19	94.29	100.00	90.75	85.98
Population (% of regional population represented in the database)						
World	81.10	92.31	91.88	93.57	90.57	89.89
<i>Milanovic (2012): World</i>	<i>87</i>	<i>92</i>	<i>92</i>	<i>94</i>		
Mature economies	94.98	99.85	99.59	96.67	96.96	97.61
China	100.00	100.00	100.00	100.00	100.00	100.00
India	100.00	100.00	100.00	100.00	100.00	100.00
Other Asia	74.64	85.42	88.70	88.71	88.70	85.23
M. East & N. Africa	60.42	69.48	63.60	68.42	47.82	61.95
Sub-Saharan Africa	28.51	72.90	68.03	80.15	74.12	64.74
L. America & Caribbean	88.20	92.89	94.91	96.41	94.43	93.37
Russia, C. Asia, SE Europe	28.42	87.41	87.41	99.40	84.41	77.41
<i>Notes: Observations are weighted by population size in the computation of global coverage, otherwise unweighted. The last column is the (unweighted) average over the 5 benchmark year values. Milanovic (2012) refers to Table 2 of that paper, using the 2002 figures for the 2003 benchmark year.</i>						

Table 4.2: Panel summary statistics: Number of countries by panel duration

Regions	Total	# benchmark years					# countries observed	
		1	2	3	4	5	in 1988 & 2008	in 1993 & 2008
World	162	22	24	27	31	58	63	90
Mature economies	39	0	2	2	10	25	29	34
China	2	0	0	0	0	2	2	2
India	2	0	0	0	0	2	2	2
Other Asia	19	2	5	1	3	8	8	11
M. East & N. Africa	11	4	1	1	3	2	2	3
Sub-Saharan Africa	43	11	9	14	5	4	4	16
L. America & Caribbean	26	4	3	0	5	14	15	17
Russia, C. Asia, SE Europe	20	1	4	9	5	1	1	5

Notes: The last two columns allow for gaps in the panel.

Figure 4.1: Comparison of global and within-country inequality




Table 4.3: Global and regional inequality

	Benchmark year									1988-2008
	1988	change (%, pp)	1993	change (%, pp)	1998	change (%, pp)	2003	change (%, pp)	2008	change (%,pp)
A. Global inequality (%)										
Gini index	72.15	-0.33	71.92	-0.52	71.54	0.45	71.86	-1.89	70.50	-2.29
GE(0) (Theil-L)	114.22	-3.06	110.72	-3.27	107.10	0.42	107.55	-4.50	102.71	-10.07
GE(1) (Theil-T)	102.20	0.24	102.45	0.34	102.80	2.09	104.94	-4.41	100.31	-1.85
GE(2)	173.71	3.14	179.16	7.75	193.04	5.82	204.28	-1.42	201.38	15.92
Atkinson A(2)	83.53	-0.83	82.83	-1.28	81.78	0.23	81.96	0.00	81.96	-1.87
Atkinson A(1)	68.09	-1.67	66.95	-1.82	65.73	0.23	65.89	-2.57	64.20	-5.71
Atkinson A(0.5)	43.53	-1.18	43.02	-1.39	42.42	0.87	42.79	-4.10	41.03	-5.74
B. Regional Gini indices (%)										
Mature economies	38.22	1.66	38.85	0.72	39.13	-0.89	38.79	8.11	41.93	9.72
China	32.02	10.70	35.45	8.63	38.51	8.48	41.78	2.32	42.74	33.47
India	31.15	-3.28	30.13	4.29	31.42	3.26	32.44	2.03	33.10	6.27
Other Asia	44.47	-0.47	44.26	5.33	46.62	-10.33	41.80	7.56	44.96	1.12
M.East & N.Africa	41.79	0.41	41.96	3.69	43.51	-9.45	39.40	-5.81	37.11	-11.21
Sub-Saharan Africa	40.56	31.93	53.51	-2.54	52.15	8.27	56.46	3.25	58.30	43.74
L.America & Carib.	52.66	3.72	54.61	3.38	56.46	-1.35	55.70	-5.21	52.79	0.26
Russia,C.Asia,SE Eu.	46.50	3.87	48.30	-16.95	40.11	4.26	41.82	0.11	41.87	-9.96
C. Country decomposition of GE(0): between-country contribution in % (change in percentage points)										
GE(0) betw. contrib.	83.19	-3.14	80.05	-1.81	78.25	-0.36	77.88	-1.20	76.68	-6.51
D. Country decomposition of GE(0): level of within-country inequality										
GE(0) (Theil-L) (%)	19.20	15.06	22.09	5.49	23.30	2.10	23.79	0.69	23.95	24.78
E. Country decomp. of GE(0): within & betw. contributions of CHN, IND & USA in % (change in pp)										
China: Between	45.37	-7.88	37.48	-4.76	32.72	-5.26	27.47	-5.55	21.92	-23.45
India: Between	32.11	-3.06	29.05	0.38	29.44	0.66	30.09	3.03	33.12	1.01
USA: Between	-8.29	0.93	-7.36	-0.51	-7.86	0.12	-7.75	-0.65	-8.40	-0.11
Rest: Between	14.01	6.87	20.88	3.07	23.95	4.12	28.07	1.97	30.04	16.03
China: Within	3.09	-0.01	3.07	0.74	3.81	0.31	4.12	0.55	4.67	1.58
India: Within	2.69	-0.41	2.28	0.45	2.73	0.09	2.81	0.35	3.16	0.47
USA: Within	1.33	-0.15	1.18	0.16	1.34	-0.04	1.30	0.30	1.59	0.26
Rest: Within	9.70	3.72	13.42	0.46	13.88	0.01	13.89	0.01	13.89	4.20
F. Absolute measures of inequality										
Absolute Gini index	2377	-0.56	2364	5.03	2483	5.09	2609	10.68	2888	21.49
G. Average annual incomes per capita (in 2005 PPP-adjusted USD), by percentile groups										
Bottom 10%	201	1.24	203	6.61	217	4.90	228	10.29	251	24.88
P40-P50	552	12.36	620	15.36	715	7.10	766	22.90	941	70.62
P50-P60	791	10.91	877	11.17	975	7.14	1045	30.01	1359	71.74
P60-P70	1323	2.26	1353	13.66	1538	5.13	1616	29.26	2089	57.95
P80-P90	7414	-3.45	7158	0.26	7177	-1.11	7097	9.26	7754	4.59
P90-P95	12960	1.47	13150	2.45	13472	5.56	14221	6.27	15113	16.61
P95-P99	21161	1.38	21452	5.63	22660	8.01	24474	9.68	26844	26.85
Top 1%	38964	1.64	39601	17.63	46583	10.86	51641	24.35	64213	64.80
H. Average annual incomes per capita (in 2005 PPP-adjusted USD), by region										
World	3295	-0.23	3287	5.59	3471	4.62	3631	12.82	4097	24.34
Mature economies	11457	7.11	12272	8.91	13366	12.37	15019	5.41	15832	38.19
China	484	18.18	572	37.94	789	29.02	1018	56.39	1592	228.93
India	538	4.09	560	13.93	638	0.63	642	12.62	723	34.39
Other Asia	671	19.82	804	9.70	882	6.92	943	19.72	1129	68.26
M.East & N.Africa	1773	5.75	1875	5.28	1974	-10.74	1762	-16.52	1471	-17.03
Sub-Saharan Africa	607	22.24	742	-3.10	719	8.34	779	-2.18	762	25.54
L.America & Carib.	3153	-5.42	2982	6.91	3188	-5.14	3024	29.00	3901	23.72
Russia,C.Asia,SE Eu.	2498	10.37	2757	-16.65	2298	10.70	2544	75.47	4464	78.70
Notes: For country decomposition, changes between benchmark years are in percentage points. For all other rows, changes are measured in percent (not annualised). Observations are weighted using population.										

Figure 4.2: Lorenz curves in 1988 and 2008





Table 4.4: Robustness check on global Gini index: Accounting for top incomes

	Benchmark year									1988- 2008 change
	1988	change (%)	1993	change (%)	1998	change (%)	2003	change (%)	2008	
1 Survey data only	72.45	-0.95	71.76	0.17	71.89	0.08	71.94	-3.30	69.57	-3.98
2 Private consumption mean	71.84	-1.55	70.72	-0.02	70.70	0.22	70.86	-3.21	68.58	-4.53
3 Censored private consump.	71.50	-1.46	70.45	0.14	70.55	0.28	70.75	-4.42	67.62	-5.43
4 Censored private consump. with Pareto imputation	71.78	-1.30	70.84	0.17	70.96	0.22	71.12	-4.31	68.05	-5.19
5 Censored private consump. with top-heavy Pareto imp.	76.34	-0.30	76.10	1.45	77.21	1.12	78.08	-2.83	75.87	-0.61
6 (5) & Pareto $a=[1.29,3.62]$	76.12	-0.42	75.80	1.16	76.68	1.09	77.51	-3.10	75.11	-1.32
7 (5) & Pareto $a=[1.5,3]$	75.56	-0.44	75.23	1.27	76.19	-0.41	75.88	-2.13	74.26	-1.72
8 (5) & share $s_{10} \leq 50\%$	75.80	-0.69	75.28	1.23	76.21	0.54	76.62	-2.75	74.51	-1.70
9 (5) & share $s_{10} \leq 40\%$	75.14	-1.02	74.37	0.47	74.72	-0.47	74.37	-2.52	72.49	-3.52
Number of surveys	63		105		112		126		114	


Notes: Observations are weighted using fractile population. All calculations are done across the sample of 520 country-years for which private consumption from national accounts is available. See description in the text: Row 5 allocates the entire gap between (censored) private consumption and survey mean to the top 10%. Censored private consumption equals the maximum of survey mean and private consumption in national accounts.

Figure 4.3: Global growth incidence curve (for fragile groups)


(a) 1988-2008



(b) For different time periods



(c) Mature economies, Asia, and Latin America & Caribbean



(d) Absolute gains




Figure 4.4: The global distribution of income




Figure 4.6: The regional composition of the global distribution of income




Figure 4.5: The changing global distribution of income




Figure 4.7: Comparing different Gini coefficients




Figure 4.8: Baseline and top-heavy distribution of income




Table 4.5: 20 biggest gainers and losers: 1988-2008

Average annualised growth		Annual growth between 1988 and 2008 contribution				
(1) decile group	(2) growth	(3) decile group	(4) growth	(5) growth	(6) share	(7) residual
El Salvador-1	9.60%	El Salvador-1	8.02%	6.55%	75.31%	18.14%
Costa Rica-1	7.95%	China-urban-10	7.73%	68.31%	9.47%	22.22%
China-urban-10	7.74%	Costa Rica-1	7.67%	25.89%	39.48%	34.63%
Ireland-1	7.23%	Ireland-1	7.16%	47.79%	21.51%	30.70%
Mexico-10	6.59%	China-urban-9	6.45%	94.29%	1.70%	4.00%
China-urban-9	6.46%	China-urban-8	6.14%	102.47%	-0.74%	-1.73%
China-urban-8	6.14%	Mexico-10	5.93%	16.52%	61.47%	22.01%
China-urban-7	5.93%	China-urban-7	5.92%	108.68%	-2.59%	-6.08%
China-urban-6	5.73%	Guatemala-1	5.82%	47.44%	26.34%	26.21%
China-rural-10	5.63%	China-urban-6	5.72%	115.03%	-4.49%	-10.54%
Ireland-2	5.57%	China-rural-10	5.60%	74.06%	10.53%	15.41%
China-urban-5	5.55%	China-urban-5	5.54%	120.93%	-6.25%	-14.68%
UK-1	5.50%	Ireland-2	5.48%	74.81%	10.37%	14.81%
China-urban-4	5.35%	UK-1	5.48%	62.58%	17.06%	20.36%
Ireland-3	5.31%	China-urban-4	5.34%	128.28%	-8.45%	-19.83%
China-rural-9	5.14%	Ireland-3	5.25%	80.13%	8.18%	11.69%
China-urban-3	5.13%	Guatemala-2	5.17%	57.19%	21.46%	21.35%
Ireland-4	5.10%	China-urban-3	5.12%	136.84%	-11.01%	-25.83%
Chile-1	5.00%	China-rural-9	5.09%	86.19%	5.61%	8.20%
China-urban-2	4.90%	Ireland-4	5.04%	85.39%	6.02%	8.59%
...	...	...	...	...	...	...
Bulgaria-1	-4.45%	Bulgaria-9	-2.93%	104.89%	-9.23%	4.35%
Lithuania-6	-4.54%	Bulgaria-8	-2.98%	103.72%	-7.02%	3.30%
Romania-8	-4.59%	Japan-1	-3.05%	13.74%	92.10%	-5.84%
Lithuania-4	-4.69%	Bulgaria-7	-3.05%	101.83%	-3.45%	1.62%
Estonia-2	-4.72%	Honduras-1	-3.13%	-254.07%	161.26%	192.80%
Romania-7	-4.90%	Romania-5	-3.13%	81.99%	29.33%	-11.32%
Estonia-3	-4.93%	Latvia-2	-3.15%	47.33%	67.83%	-15.17%
Lithuania-5	-4.95%	Paraguay-1	-3.16%	14.88%	91.58%	-6.46%
Estonia-7	-5.07%	Bulgaria-6	-3.16%	99.31%	1.30%	-0.61%
Romania-6	-5.15%	Bulgaria-5	-3.28%	96.75%	6.13%	-2.89%
Romania-5	-5.33%	Bulgaria-4	-3.41%	94.03%	11.27%	-5.31%
Estonia-4	-5.34%	Bulgaria-3	-3.54%	91.65%	15.77%	-7.42%
Romania-4	-5.48%	Romania-4	-3.54%	75.07%	40.60%	-15.67%
Estonia-6	-5.55%	Bulgaria-2	-3.71%	88.65%	21.45%	-10.09%
Romania-3	-5.60%	Romania-3	-4.11%	67.92%	52.24%	-20.16%
Estonia-5	-5.60%	Latvia-1	-4.27%	38.40%	79.34%	-17.74%
Romania-2	-5.75%	Romania-2	-4.61%	63.15%	60.01%	-23.16%
Bolivia-1	-6.00%	Bulgaria-1	-4.71%	76.00%	45.34%	-21.34%
Estonia-1	-6.36%	Romania-1	-6.15%	53.67%	75.45%	-29.12%
Romania-1	-7.45%	Bolivia-1	-6.67%	-39.19%	107.63%	31.57%

Notes: Only for countries observed in 1988 and 2008 (might include gaps, 63 countries); columns 1 & 2 exclude countries with no 5-year intervals (Cyprus). Decile groups are numbered 1 to 10, with 1 being the bottom decile. The decomposition is based on the 20-year growth rate, but here the annualised growth rates are displayed. A positive contribution implies that source (growth/share/residual) works in same direction as growth rate.

Figure 4.9: Quasi-non-anonymous growth incidence curve





Figure 4.10: Positional gains



4.A Appendix

In this section, we (1) report the size of the national accounts adjustment, and (2) compare the Pareto coefficients and top shares to those in the WTID. To summarise the magnitude of our national accounts adjustment, Table 4.B.3 presents the minimum and average values of the ratio “survey mean to mean used in the top-heavy Pareto adjustment” across the five benchmark years and across the different regions.¹²⁶ Across the five benchmark years, the survey mean is on average between 75.15% and 81.43% of (censored) national accounts consumption.¹²⁷ The national accounts adjustment can, however, be substantially greater. The minimum value over all surveys is 22.2% (Swaziland in 2008). The low values for the mature economies are due to the new EU member states.

We also match the surveys in our database with information from the WTID for the same survey year. In Tables 4.B.4 and 4.B.5 we compare the Pareto constants and the top shares respectively. For the country-years we could match, the Pareto coefficient in the WTID ranges from 1.29 to 3.62 across the years (panel C in Table 4.B.4). As explained in Section 4.4, we use two Pareto imputations; the “proportional” and “top-heavy” adjustments. The proportional adjustment tends to produce Pareto coefficients that are too high compared with the WTID, thus understating inequality (compare panels A and C in Table 4.B.4). This is not surprising given that we expect household surveys, from which the decile group shares in this imputation are taken, to underreport top incomes. The Pareto coefficients obtained using the “top-heavy” adjustment seem closer to the values in the WTID (compare panels B and C in Table 4.B.4). However, the lowest values in this adjustment (just above 1) are implausibly low, thus overstating inequality. This is the result of some large national accounts discrepancies allocated entirely to the top 10%. Furthermore, only around 70% of the

¹²⁶In other words, Table 4.B.3 shows the ratio of survey mean to $\max(\text{survey mean, private national accounts consumption})$, which by definition is less than or equal to 1.

¹²⁷Excluding observations for which the survey mean exceeds national accounts consumption (i.e. the ratio in Table 4.B.3 equals 100%), brings this average to between 67.44% and 75.62% across the five benchmark years. The proportion of observations for which the survey mean exceeds national accounts consumption ranges from 18.3% to 26.7%.

Pareto constants are within the range observed in the WTID, compared with 90% in the proportional adjustment.

We can also compare our top income shares with those from the WTID (Table 4.B.5).¹²⁸ For each benchmark year and three different top shares (10%, 5% and 1%) we provide summary statistics on our own shares using the top-heavy adjustment, the WTID shares and the difference between the two. The top 1% share is typically too high. There are some very high values of about 70%, which are clearly excessive, but arise out of the allocation of the national accounts excess to the top 10%. The P95-P99 shares are relatively similar to the WTID and the P90-P95 shares are too small. It seems that this is the consequence of the Pareto imputation applied to the excessive top 10% share which has resulted in the tails being too fat.

¹²⁸We do not repeat the same exercise for the proportional Pareto adjustment because the information contained in the shares is also in the Pareto constant, which we compare in Table 4.B.4.

4.B Appendix Tables

Table 4.B.1: Comparison with previous estimates of global inequality

	Benchmark years							
	1988	change (%, pp)	1993	change (%, pp)	1998	change (%, pp)	2003	2008
	Global inequality (change in percent)							
Gini index (%)								
<i>Own estimates (Table 4.3)</i>	72.15	-0.33	71.92	-0.52	71.54	0.45	71.86	70.50
Milanovic (2012)	67.8	2.2	69.3	-0.7	68.8	1.9	70.1	
Milanovic (2005)	61.9	5.3	65.2	-1.5	64.2			
Milanovic (2002)	62.5	5.4	65.9					
Bourguignon (2012)	71				69			66
Warner et al (2012)			70.0	-1.4	69.0	-2.9	67.0	
Chotikapanich et al. (2012)			64.8	-1.2	64.0			
Bhalla (2002) (Income)	67				65			
Bhalla (2002) (Consumption)	66				63			
Bourguignon & Morrisson (2002)			66					
Chotikapanich et al. (1997)	65							
Dikhanov & Ward (2002)	69				68			
Dowrick & Akmal (2005) (GK)			64					
Dowrick & Akmal (2005) (Afriat)			71					
Sala-i-Martin (2006)	65	-1.4	64	-0.3	64			
GE(0) (Theil-L or mean-log-deviation) (%)								
<i>Own estimates (Table 4.3)</i>	114.22	-3.06	110.72	-3.27	107.10	0.42	107.55	102.71
Milanovic (2002)	75.8	13.98	86.4					
Warner et al (2012)			105.3	-6.3	98.6	-11.1	87.7	
Chotikapanich et al. (2012)			81.3	-2.2	79.5			
Chotikapanich et al. (1997)	80.6							
Dikhanov & Ward (2002)	102.1				97.1			
Sala-i-Martin (2006)	84.2	-2.7	81.9	-0.4	81.6			
GE(1) (Theil-T) (%)								
<i>Own estimates (Table 4.3)</i>	102.20	0.24	102.45	0.34	102.80	2.09	104.94	100.31
Milanovic (2012)	87.5	7.09	93.7	0.53	94.2	5.94	99.8	
Milanovic (2005)	71.5	14.41	81.8	-3.18	79.2			
Bourguignon & Morrisson (2002)			85.5					
Dikhanov & Ward (2002)	89.1				90.7			
Dowrick & Akmal (2005) (GK)			79					
Dowrick & Akmal (2005) (Afriat)			101					
Sala-i-Martin (2006)	80.8	-2.6	78.7	-0.3	78.5			
Decomposition by country (change in percentage points)								
GE(0) between-country contribution (%)								
<i>Own estimates (Table 4.3)</i>	83.19	-3.14	80.05	-1.81	78.25	-0.36	77.88	76.68
Milanovic (2002) (common sample)	75	-1	74					
Warner et al (2012)			73	-3	70	-5	65	
Chotikapanich et al. (2012)			65	-3	62			
Sala-i-Martin (2006)	68	-3	65	-3	62			
<i>Notes: Milanovic (2012): Table 4, p. 14: Gini from row 5 (2005 PPP, sep. rural-urban prices for China, India & Indonesia); Theil from row 3 (2005 PPP, sep. rural-urban prices for China only); 2002 figures for 2003 benchmark. Milanovic (2002): Table 16, p. 72: Using full sample; Table 19, p. 78 (decomposition): Only for common sample. Milanovic (2005): Table 9.4, p. 108: Using full sample. Bourguignon (2012): Figure 1, only approximate, because read-off from figure; 1988 refers to 1989, 1998 refers to 1997, 2008 refers to 2006. Warner et al (2012): Table 4, p. 31; 1998 refers to 2000, 2003 refers to 2005. Chotikapanich et al. (2012): Table 4, p. 66; 1998 refers to 2000. Otherwise: Anand and Segal (2008), Table 1: Survey estimates allocated to benchmark according to rules with micro data: 1988: Bhalla, Chotikapanich et al, and Dikhanov & Ward refer to 1990; 1993: Bourguignon & Morrisson refer to 1992; 1998: Bhalla refers to 2000, and Dikhanov & Ward refer to 1999.</i>								

Table 4.B.2: Global Gini index (%) for balanced sample

	Benchmark year								
	1988	change (%)	1993	change (%)	1998	change (%)	2003	change (%)	2008
Baseline (Table 3)	72.15	-0.33	71.92	-0.52	71.54	0.45	71.86	-1.89	70.50
<i>N</i>	(75)		(115)		(121)		(133)		(121)
Balanced ('88-'08)	72.76	0.07	72.82	-0.70	72.31	0.26	72.50	-2.25	70.87
<i>N</i>	(58)		(58)		(58)		(58)		(58)
Balanced ('93-'08)			72.00	-0.36	71.74	0.00	71.74	-2.41	70.01
<i>N</i>			(78)		(78)		(78)		(78)
Balanced ('98-'08)					71.82	-0.01	71.81	-2.38	70.10
<i>N</i>					(92)		(92)		(92)
Balanced ('03-'08)							72.06	-2.27	70.42
<i>N</i>							(107)		(107)
<i>Notes: Table shows global Gini index (%) for the unbalanced baseline and four balanced samples (depending on start year). Number of countries in parentheses.</i>									

Table 4.B.3: Size of national accounts adjustment: Survey mean to max(survey mean, national accounts consumption) (%)

	Benchmark years									
	1988		1993		1998		2003		2008	
	mean	min	mean	min	mean	min	mean	min	mean	min
World	81.4	41.3	80.5	25.4	80.8	22.5	78.8	36.9	75.1	22.2
Mature economies	84.6	41.3	84.0	29.4	80.5	55.7	78.9	42.9	75.3	35.8
China	100.0	100.0	90.6	90.6	95.3	95.3	92.2	92.2	100.0	100.0
India	77.5	77.5	71.7	71.7	60.0	60.0	56.5	56.5	53.0	53.0
Other Asia	67.1	43.6	72.0	41.5	78.1	34.9	73.9	48.3	74.2	50.8
M. East & N. Africa	80.2	61.2	84.4	54.3	85.3	57.4	80.3	54.1	67.5	46.7
Sub-Saharan Africa	91.1	52.6	80.0	25.6	84.5	22.5	87.1	53.2	81.4	22.2
L. America & Carib.	80.1	48.5	75.7	25.4	81.3	43.4	73.0	45.2	73.3	45.6
Russia,C.Asia,SE Eu.	58.2	58.2	84.9	43.5	76.2	49.1	74.4	36.9	69.5	36.3

Notes: Observations are unweighted. Only for country-decile-groups which have national accounts final household consumption. By definition, the values in the table are $\leq 100\%$.

Table 4.B.4: Comparison of Pareto constant with fiscal data

	Benchmark years				
	1988	1993	1998	2003	2008
A. Proportional Pareto imputation (decile group shares are unchanged)					
Mean	3.02	2.58	2.57	2.54	2.61
Median	2.66	2.47	2.40	2.41	2.56
Min	1.65	1.37	1.66	1.43	1.27
Max	12.78	4.58	4.14	4.19	4.29
Percentage within [1.29, 3.62] (WTID range)	80.95	88.57	90.18	92.86	92.98
Number of surveys	63	105	112	126	114
B. Top-heavy Pareto: national accounts excess to top 10%					
Mean	1.99	1.85	1.77	1.71	1.69
Median	1.64	1.65	1.65	1.54	1.43
Min	1.13	1.06	1.06	1.12	1.06
Max	4.58	4.58	4.14	3.89	3.39
Percentage within [1.29, 3.62] (WTID range)	74.60	77.14	77.68	77.78	68.42
Number of surveys	63	105	112	126	114
C. World Top Incomes Database					
Mean	2.46	2.39	2.06	1.95	1.97
Median	2.43	2.50	1.87	2.00	2.01
Min	1.53	1.37	1.29	1.48	1.56
Max	3.20	3.19	3.62	2.61	2.33
Number of surveys	14	18	21	18	12
<i>Notes: Observations are unweighted. Only for country-decile-groups which have national accounts final household consumption. World Top Incomes Database: http://topincomes.g-mond.parisschoolofeconomics.eu/, accessed 24.01.2013.</i>					

Table 4.B.5: Comparing top-heavy shares with the World Top Incomes Database

	Benchmark years														
	1988			1993			1998			2003			2008		
	own	WTID	diff	own	WTID	diff	own	WTID	diff	own	WTID	diff	own	WTID	diff
Top 1 % share of total income															
mean	19.7	7.4	9.3	22.7	8.1	8.5	21.6	10.1	8.9	23.2	11.5	8.7	25.7	11.5	5.9
median	16.6	7.6	6.4	17.6	7.9	4.6	17.1	10.2	6.1	20.6	9.8	5.1	24.5	9.7	4.4
min	3.2	4.9	-4.6	3.0	5.0	-5.1	3.9	5.2	-4.6	3.8	5.6	-4.6	4.9	6.1	-9.4
max	53.2	10.5	40.7	75.9	12.2	43.4	76.3	18.1	46.2	55.9	19.9	33.7	75.5	20.3	28.7
N	63	16	16	105	20	20	112	23	23	126	19	19	114	15	15
P95-P99 share of total income															
mean	13.5	12.5	-0.3	13.7	13.0	-0.6	14.0	13.0	0.0	14.1	13.9	-0.7	13.6	14.2	-1.2
median	13.9	13.1	-0.5	14.2	13.1	-0.4	14.4	13.0	0.6	14.4	13.3	-0.2	13.6	13.2	-0.1
min	8.1	10.1	-4.4	7.6	10.2	-4.7	7.1	10.2	-5.8	8.7	10.4	-6.7	7.6	10.6	-5.8
max	18.1	14.8	3.7	18.1	18.7	2.4	18.2	15.5	2.3	17.6	20.3	2.4	17.5	22.0	1.5
N	63	13	13	105	17	17	112	18	18	126	16	16	114	13	13
P90-P95 share of total income															
mean	9.2	11.0	-1.9	9.2	10.9	-1.5	9.3	10.8	-1.5	9.2	11.0	-1.8	8.7	11.0	-1.7
median	9.3	10.9	-1.8	9.4	10.9	-1.6	9.4	10.9	-1.6	9.3	10.7	-1.5	9.0	10.9	-1.6
min	5.3	9.2	-4.1	3.6	8.8	-3.9	3.3	8.5	-4.7	5.1	9.1	-6.0	3.5	9.0	-5.6
max	11.9	12.8	-0.2	12.8	13.3	1.6	12.9	13.3	2.3	11.9	14.6	0.7	12.1	14.6	0.3
N	63	12	12	105	16	16	112	18	18	126	15	15	114	12	12
<i>Notes: Observations are unweighted. Only for countries which have national accounts final household consumption. "own" refers to the shares obtained from the top-heavy Pareto imputation. WTID: accessed 24.01.2013. "diff" is defined as the difference between the top-heavy and WTID shares; a positive difference implies that the share from the top-heavy Pareto imputation exceeds the share observed in the WTID.</i>															

Table 4.B.6: 20 biggest gainers and losers: 1993-2008

Average annualised growth		Annual growth between 1993 and 2008				
(1) decile group	(2) growth	(3) decile group	(4) growth	contribution		
				(5) growth	(6) share	(7) residual
Swaziland-1	19.09%	Panama-1	17.53%	2.44%	78.03%	19.53%
Panama-1	17.68%	Lithuania-9	9.56%	92.67%	1.97%	5.36%
Swaziland-2	16.31%	Lithuania-8	9.52%	93.35%	1.79%	4.86%
Kenya-1	16.17%	Lithuania-7	9.51%	93.53%	1.74%	4.73%
Lithuania-1	15.86%	Lithuania-6	9.38%	95.76%	1.14%	3.10%
Romania-10	15.36%	Lithuania-4	9.33%	96.66%	0.90%	2.44%
Swaziland-3	14.19%	Lithuania-5	9.27%	97.87%	0.57%	1.56%
Azerbaijan-2	13.93%	Lithuania-3	9.19%	99.27%	0.20%	0.54%
Azerbaijan-1	13.79%	Swaziland-1	9.11%	48.97%	21.98%	29.04%
Azerbaijan-3	13.79%	Centr. Afr. Rep.-1	8.86%	41.64%	28.19%	30.18%
Azerbaijan-4	13.71%	Lithuania-10	8.69%	109.24%	-2.48%	-6.75%
Romania-9	13.54%	Lithuania-2	8.65%	109.94%	-2.67%	-7.27%
Azerbaijan-5	13.45%	Centr. Afr. Rep.-2	8.47%	44.90%	26.61%	28.49%
Azerbaijan-6	13.24%	Swaziland-2	8.22%	58.16%	18.02%	23.81%
Romania-8	13.08%	China-urban-10	8.14%	75.45%	9.13%	15.41%
Azerbaijan-7	12.97%	Lithuania-1	7.95%	126.53%	-7.13%	-19.39%
Swaziland-4	12.96%	Centr. Afr. Rep.-3	7.72%	52.23%	23.07%	24.70%
Romania-7	12.83%	Swaziland-3	7.67%	65.10%	15.04%	19.86%
Lithuania-2	12.73%	Latvia-10	7.64%	70.11%	12.38%	17.50%
Romania-6	12.67%	Ireland-1	7.58%	60.15%	18.13%	21.72%
...	...	...	...	...	...	...
Japan-3	-2.09%	Kyrgyz Rep.-10	-3.93%	49.38%	65.17%	-14.54%
Burundi-7	-2.11%	Tanzania-6	-3.96%	102.96%	-5.57%	2.61%
Slovak Rep.-1	-2.13%	Tanzania-5	-4.02%	101.69%	-3.18%	1.49%
Bulgaria-1	-2.24%	Tanzania-10	-4.07%	100.92%	-1.73%	0.81%
Israel-1	-2.42%	Tanzania-8	-4.09%	100.42%	-0.79%	0.37%
Centr. Afr. Rep-3	-2.53%	Tanzania-7	-4.15%	99.49%	0.96%	-0.45%
Burundi-6	-2.54%	Tanzania-4	-4.15%	99.44%	1.04%	-0.49%
Kyrgyz Rep.-10	-2.72%	Kenya-5	-4.19%	145.87%	-148.34%	102.47%
Japan-2	-2.95%	Tanzania-9	-4.20%	98.56%	2.71%	-1.27%
Burundi-5	-3.15%	Tanzania-1	-4.25%	97.65%	4.41%	-2.06%
Centr. Afr. Rep-2	-3.29%	Tanzania-3	-4.27%	97.43%	4.83%	-2.26%
Burundi-4	-3.90%	Tanzania-2	-4.33%	96.37%	6.83%	-3.19%
Japan-1	-4.52%	Kenya-6	-4.74%	133.54%	-108.48%	74.94%
Burundi-3	-5.12%	Japan-1	-4.79%	35.37%	79.24%	-14.62%
Bolivia-1	-5.30%	Kenya-7	-5.35%	123.02%	-74.46%	51.43%
Kenya-10	-5.79%	Kenya-8	-6.00%	114.19%	-45.90%	31.71%
Honduras-1	-6.31%	Bolivia-1	-6.18%	-9.10%	103.31%	5.79%
Centr. Afr. Rep-1	-6.51%	Honduras-1	-6.65%	-108.94%	122.79%	86.15%
Burundi-2	-7.34%	Kenya-9	-7.04%	103.79%	-12.25%	8.46%
Burundi-1	-13.11%	Kenya-10	-11.94%	81.12%	61.06%	-42.18%

Notes: Only for countries observed in 1993 and 2008 (might include gaps, 90 countries); columns 1 & 2 exclude countries with no 5-year intervals (Niger). Deciles are numbered 1 to 10, with 1 being the bottom decile. The decomposition is based on the 20-year growth rate, but here the annualised growth rates are displayed. A positive contribution implies that source (growth/share/residual) works in same direction as growth rate.

Conclusion

In this thesis I investigated four different aspects of income inequality at the national and global level. I have worked with very different types of data - price data, administrative social security and tax records, surveys of employers and household surveys - in order to deepen our understanding of inequality and to highlight the subtleties of what is meant by this term.

In the first chapter, I showed that regional prices have an equalising effect on district wage differentials in Germany. This effect is stronger in West than East Germany because the price distribution in the East is more concentrated than in the West relative to the dispersion of average wages. Furthermore, I found that the change in aggregate inequality as a result of controlling for regional price differences is small, although significant. Between-district inequality falls by 40%. Consistent with the results for district averages, this effect is weaker in East (around 35%) but stronger in West Germany (almost 60%). The effect on aggregate inequality is small because aggregate inequality is mostly explained by differences within rather than between districts.

The second chapter contributes to the recent literature on top income shares in the United States. I found that the pattern of top income shares is robust to controlling for tax unit size. This suggests that the fall in average tax unit size between the 1960s and 2000s was relatively uniform across the distribution. In the second part of the chapter, I analyse the joint distribution of capital and labour income. Over the same period as top income shares increased, the composition of these incomes changed dramatically, with the labour share rising. Income is decomposed into two sources - labour and capital, with business income split between the two sources. The marginal distributions of

labour and capital income have both become more unequal. I used a copula framework to show that these factor incomes have become more closely associated at least at the top. This can explain why the top income share in terms of total income increased faster than in either of the marginal distributions. This association is asymmetric in the following sense: Those tax units that are at the top of the wage distribution are more likely to also be at the top of the distribution of capital incomes, compared with those tax units at the top of the capital distribution being at the top of the distribution of wages. Hence many of the richest Americans are highly paid and also receive high incomes from capital, while there remain some rentiers who receive little income from work.

In the third chapter, I estimated the enterprise size effect in wage- and self-employment in urban Ghana. I found that individuals in larger enterprises earn higher wages in the cross-section, even with controls for observable human capital. In wage-employment, the size premium is robust to controlling for time-invariant unobservables. On the other hand, the size premium for the self-employed seems to be driven by time-invariant unobservables which is robust to using dynamic panel estimators.

The final chapter moved from the national to the global level and investigated interpersonal inequality around the world between 1988 and 2008. The chapter put together a database of national household surveys which improved on the datasets used in the earlier literature. The global Gini index remains high at around 70.5% in 2008, having declined by approximately 2pp over this 20-year period. Thanks to rapid growth of average income particularly in China, between-country inequality declined, while within-country inequality increased. I mostly considered relative measures of inequality so it is important to bear in mind that 44% of the increase of global income went to the top 5% of world population. When I adjusted for the likely underreporting of top incomes in surveys, I estimated a higher global Gini of 75.9% in 2008 and no decline over time. Tracking the evolution of individual country-decile-groups shows that China graduated from the bottom ranks, thereby transforming a twin-peaked 1988 global distribution into a single-peaked one now. As a result of the growth in

Asia, the poorest segments of the global distribution are now largely from Sub-Saharan Africa. The “winners” (fastest growth) were country-deciles that in 1988 were around the median of the global income distribution, 90% of whom in terms of population are from Asia. The “losers” (slowest growth) were the country-deciles that in 1988 were around the 85th percentile of the global income distribution, almost 90% of whom in terms of population are from mature economies.

Bibliography

- Aaberge, R. and A. B. Atkinson: 2010, 'Top incomes in Norway'. In: A. B. Atkinson and T. Piketty (eds.): *Top Incomes: A Global Perspective*. Oxford University Press.
- Aaberge, R., A. B. Atkinson, S. Königs, and C. Lakner: 2013, 'From Classes to Copulas: Wages, capital and top incomes'. Working paper, University of Oxford.
- Aaberge, R., M. Bhuller, A. Langørgen, and M. Mogstad: 2010, 'The distributional impact of public services when needs differ'. *Journal of Public Economics* **94**(9), 549–562.
- Abdelkrim, A. and J.-Y. Duclos: 2007, 'DASP: Distributive Analysis Stata Package'. PEP, World Bank, UNDP and Université Laval.
- Abowd, J. M. and F. Kramarz: 2000, 'Inter-industry and Firm-size Wage Differentials: New Evidence from Linked Employer-Employee Data'. Working paper, Cornell University.
- Abowd, J. M., F. Kramarz, and D. N. Margolis: 1999, 'High Wage Workers and High Wage Firms'. *Econometrica* **67**(2), 251–333.
- Ackland, R., S. Dowrick, and B. Freyens: 2013, 'Measuring Global Poverty: Why PPP Methods Matter'. *Review of Economics and Statistics* **95**(3), 813–824.
- Alessandria, G. and J. P. Kaboski: 2011, 'Pricing-to-Market and the Failure of Absolute PPP'. *American Economic Journal: Macroeconomics* **3**(1), 91–127.
- Alvaredo, F.: 2010, 'The Rich in Argentina over the Twentieth Century, 1932-2004'. In: A. B. Atkinson and T. Piketty (eds.): *Top Incomes: A Global Perspective*. Oxford University Press, pp. 253–298.
- Alvaredo, F.: 2011, 'A note on the relationship between top income shares and the Gini coefficient'. *Economics Letters* **110**(3), 274–277.
- Alvaredo, F., A. B. Atkinson, T. Piketty, and E. Saez: 2013a, 'The Top 1 Percent in International and Historical Perspective'. *Journal of Economic Perspectives* **27**(3), 3–20.
- Alvaredo, F., A. B. Atkinson, T. Piketty, and E. Saez: 2013b, 'The World Top Incomes Database'. <http://topincomes.g-mond.parisschoolofeconomics.eu/>, Accessed 24/01/2013.
- Alvaredo, F. and E. Pisano: 2010, 'Top Incomes in Italy, 1974-2004'. In: A. B. Atkinson and T. Piketty (eds.): *Top Incomes: A Global Perspective*. Oxford University Press.
- Anand, S.: 1983, *Inequality and poverty in Malaysia: measurement and decomposition*. Published for the World Bank by Oxford University Press.
- Anand, S. and P. Segal: 2008, 'What Do We Know about Global Income Inequality?'. *Journal of Economic Literature* **46**(1), 57–94.

- Anand, S., P. Segal, and J. E. Stiglitz: 2010, 'Introduction'. In: S. Anand, P. Segal, and J. E. Stiglitz (eds.): *Debates on the Measurement of Global Poverty*. Oxford University Press.
- Andrews, D. and A. C. Sánchez: 2011, 'The Evolution of Homeownership Rates in Selected OECD Countries: Demographic and Public Policy Influences'. *OECD Journal: Economic Studies* **2011**(1), 207–243.
- Arai, M.: 2003, 'Wages, Profits, and Capital Intensity: Evidence from Matched Worker-Firm Data'. *Journal of Labor Economics* **21**(3), 593–618.
- Arellano, M.: 1987, 'Computing Robust Standard Errors for Within-Group Estimators'. *Oxford Bulletin of Economics and Statistics* **49**, 431–434.
- Arellano, M.: 2003, *Panel data econometrics*. Oxford University Press, Oxford.
- Arellano, M. and S. Bond: 1991, 'Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations'. *The Review of Economic Studies* **58**(2), 277–297.
- Armour, P., R. V. Burkhauser, and J. Larrimore: 2013, 'Deconstructing Income and Income Inequality Measures: A Crosswalk from Market Income to Comprehensive Income'. *American Economic Review: Papers and Proceedings* **103**(3), 173–77.
- Ashenfelter, O.: 2012, 'Comparing Real Wage Rates: Presidential Address'. *American Economic Review* **102**(2), 617–42.
- Atkinson, A.: 2008, *The Changing Distribution of Earnings in OECD Countries*. Oxford University Press.
- Atkinson, A.: 2011, 'On lateral thinking'. *Journal of Economic Inequality* **9**, 319–328.
- Atkinson, A., A. Maynard, and C. Trinder: 1983, *Parents and children: incomes in two generations*. Heinemann Educational Books.
- Atkinson, A. B.: 1970, 'On the measurement of inequality'. *Journal of Economic Theory* **2**(3), 244–263.
- Atkinson, A. B.: 1975, *The Economics of Inequality*. Clarendon Press.
- Atkinson, A. B.: 1980, 'Horizontal equity and the distribution of the tax burden'. In: H. J. Aaron and M. J. Boskin (eds.): *The Economics of Taxation*. Washington DC: Brookings Institution.
- Atkinson, A. B.: 1981, 'The measurement of economic mobility'. In: P. J. Eijgelshoven and L. J. van Gemerden (eds.): *Inkomensverdeling en openbare financiën*. Het Spectrum, Utrecht. Reprinted in A B Atkinson, 1983, Social justice and public policy, Chapter 3, Wheatsheaf Books, Brighton.
- Atkinson, A. B.: 2003, 'Income Inequality in OECD Countries: Data and Explanations'. *CEifo Economic Studies* **49**(4), 479–513.
- Atkinson, A. B.: 2007a, 'The Distribution of Top Incomes in the United Kingdom 1908-2000'. In: A. B. Atkinson and T. Piketty (eds.): *Top Incomes over the Twentieth Century: A Contrast Between Continental European and English-Speaking Countries*. Oxford University Press.

- Atkinson, A. B.: 2007b, ‘Measuring Top Incomes: Methodological Issues’. In: A. B. Atkinson and T. Piketty (eds.): *Top Incomes over the Twentieth Century: A Contrast Between Continental European and English-Speaking Countries*. Oxford University Press.
- Atkinson, A. B.: 2009, ‘Factor shares: the principal problem of political economy?’. *Oxford Review of Economic Policy* **25**(1), 3–16.
- Atkinson, A. B.: 2013a, ‘Putting people first and macro-economic policy’. Working paper, University of Oxford.
- Atkinson, A. B.: 2013b, ‘Reducing income inequality in Europe’. *IZA Journal of European Labor Studies* **2**, 12.
- Atkinson, A. B.: 2014, ‘The colonial legacy: Income inequality in former British African colonies’. Working paper, UNU-WIDER.
- Atkinson, A. B. and A. Brandolini: 2010, ‘On Analyzing the World Distribution of Income’. *The World Bank Economic Review* **24**(1), 1–37.
- Atkinson, A. B. and J. Micklewright: 1992, *Economic transformation in Eastern Europe and the distribution of income*. Cambridge: Cambridge University Press.
- Atkinson, A. B. and T. Piketty: 2007, *Top incomes over the twentieth century: a contrast between continental European and English-speaking countries*. Oxford University Press.
- Atkinson, A. B. and T. Piketty: 2010, *Top incomes: a global perspective*. Oxford University Press.
- Atkinson, A. B., T. Piketty, and E. Saez: 2011, ‘Top Incomes in the Long Run of History’. *Journal of Economic Literature* **49**(1), 3–71.
- Atkinson, A. B., L. Rainwater, and T. M. Smeeding: 1994, ‘Income Distribution in European Countries’. LIS Working Paper Series 121, Luxembourg Income Study.
- Atkinson, A. B., L. Rainwater, and T. M. Smeeding: 1995, *Income distribution in OECD countries: evidence from the Luxembourg Income Study*. OECD social policy studies.
- Auerbach, A. J. and K. A. Hassett: 2002, ‘A New Measure of Horizontal Equity’. *The American Economic Review* **92**(4), 1116–1125.
- Auerbach, A. J. and J. Slemrod: 1997, ‘The Economic Effects of the Tax Reform Act of 1986’. *Journal of Economic Literature* **35**(2), 589–632.
- Autor, D. H., D. Dorn, and G. H. Hanson: 2013, ‘The China Syndrome: Local Labor Market Effects of Import Competition in the United States’. *American Economic Review* **103**(6), 2121–2168.
- Bach, S., G. Corneo, and V. Steiner: 2009, ‘From Bottom To Top: The Entire Income Distribution In Germany, 1992-2003’. *Review of Income and Wealth* **55**(2), 303–330.
- Baker, D.: 2013, ‘Thomas Edsall on Richard Burkhauser and Inequality’. <http://www.cepr.net/index.php/blogs/beat-the-press/thomas-edsall-on-richard-burkhauser-and-inequality>; 27 June 2013; accessed 21 February 2014.
- Baltagi, B. H., U. Blien, and K. Wolf: 2009, ‘New evidence on the dynamic wage curve for Western Germany: 1980-2004’. *Labour Economics* **16**(1), 47–51.

- Banerjee, A. and T. Piketty: 2010, 'Top Indian Incomes, 1922-2000'. In: A. B. Atkinson and T. Piketty (eds.): *Top Incomes: A Global Perspective*. Oxford University Press.
- Bardasi, E., S. Jenkins, H. Sutherland, H. Levy, and F. Zantomio: 2012, 'British Household Panel Survey Derived Current and Annual Net Household Income Variables, Waves 1-18, 1991-2009'.
- Barro, R. J. and X. Sala-i Martin: 1992, 'Convergence'. *Journal of Political Economy* **100**(2), 223–251.
- BBSR: 2009, 'Bundesinstitut für Bau-, Stadt- und Raumforschung (BBSR): Regionaler Preisindex'. *Berichte des BBR (Bundesamt für Bauwesen und Raumordnung)* **30**.
- Becker, G. S.: 1965, 'A Theory of the Allocation of Time'. *The Economic Journal* **75**(299), 493–517.
- Bell, B. D. and J. Van Reenen: 2013, 'Extreme Wage Inequality: Pay at the Very Top'. *American Economic Review: Papers and Proceedings* **103**(3), 153–57.
- Belzil, C. and J. Hansen: 2002, 'Unobserved Ability and the Return to Schooling'. *Econometrica* **70**(5), 2075–2091.
- Bhagwati, J.: 2004, *In defense of globalization*. Oxford University Press.
- Bhalla, S. S.: 2002, *Imagine There's No Country: Poverty, Inequality, and Growth in the Era of Globalization*. Institute for International Economics.
- Bivens, J. and L. Mishel: 2013, 'The Pay of Corporate Executives and Financial Professionals as Evidence of Rents in Top 1 Percent Incomes'. *Journal of Economic Perspectives* **27**(3), 57–78.
- Blanchflower, D. G.: 2000, 'Self-employment in OECD countries'. *Labour Economics* **7**(5), 471–505.
- Blanchflower, D. G. and A. J. Oswald: 2013, 'Does High Home-Ownership Impair the Labor Market?'. Working paper, Peterson Institute for International Economics.
- Blien, U., H. Gartner, H. Stüber, and K. Wolf: 2009, 'Regional price levels and the agglomeration wage differential in western Germany.'. *Annals of Regional Science* **43**(1), 71–88.
- Bloom, N., A. Mahajan, D. McKenzie, and J. Roberts: 2010, 'Why Do Firms in Developing Countries Have Low Productivity?'. *American Economic Review Papers and Proceedings* **100**(2), 619–23.
- BLS: 2013, 'Questions and Answers: Consumer Price Index Research Series Using Current Methods'. Bureau of Labor Statistics; <http://www.bls.gov/cpi/cpirsdc.htm>; Accessed 16.11.2013.
- Blundell, R. and S. Bond: 1998, 'Initial conditions and moment restrictions in dynamic panel data models'. *Journal of Econometrics* **87**, 115–143.
- Bond, S. and F. Windmeijer: 2002, 'Finite sample inference for GMM estimators in linear panel data models'. Working paper, The Institute for Fiscal Studies, London.
- Bonhomme, S. and J.-M. Robin: 2009, 'Assessing the Equalizing Force of Mobility Using Short Panels: France, 1990-2000'. *Review of Economic Studies* **76**(1), 63–92.

- Bosmans, K., K. Decancq, and A. Decoster: 2013, 'The relativity of decreasing inequality between countries'. Working paper, University of Antwerp.
- Bourguignon, F.: 1979, 'Decomposable Income Inequality Measures'. *Econometrica* **47**(4), 901–920.
- Bourguignon, F.: 2011a, 'Non-anonymous growth incidence curves, income mobility and social welfare dominance'. *The Journal of Economic Inequality* **9**(4), 605–627.
- Bourguignon, F.: 2011b, 'A Turning Point in Global Inequality ... and Beyond'. Conference presentation, ABCDE conference.
- Bourguignon, F.: 2012, *La mondialisation de l'inégalité*. Seuil et La République des Idées.
- Bourguignon, F. and C. Morrisson: 2002, 'Inequality among World Citizens: 1820-1992'. *The American Economic Review* **92**(4), 727–744.
- Brück, T. and H. Peters: 2009, '20 Years of German Unification: Evidence on Income Convergence and Heterogeneity'. IZA Discussion Paper 4454.
- Brück, T. and H. Peters: 2010, 'Persönliches Einkommen in Ostdeutschland um ein Fünftel niedriger als im Westen'. *DIW Wochenbericht* **77**, 14–20.
- Breuer, C. C. and J. Mehrhoff: 2009, 'Einkommensgruppenspezifische Teuerungsraten: Wie repräsentativ ist ein aggregiertes Preismaß?'. *Wirtschaftsdienst* **11**, 755–762.
- Brown, C. and J. Medoff: 1989, 'The Employer Size-Wage Effect'. *The Journal of Political Economy* **97**(5), 1027–1059.
- Buhmann, B., L. Rainwater, G. Schmaus, and T. M. Smeeding: 1988, 'Equivalence Scales, Well-being, Inequality, and Poverty: Sensitivity Estimates Across Ten Countries Using the Luxembourg Income Study (LIS) Database'. *Review of Income and Wealth* **34**(2), 115–142.
- Bulow, J. and L. Summers: 1986, 'A theory of dual labor markets, with applications to industrial policy, discrimination, and Keynesian unemployment'. *Journal of Labor Economics* **4**, 376–414.
- Burkhauser, R. V., S. Feng, and S. P. Jenkins: 2009, 'Using the P90/P10 Index to Measure U.S. Inequality Trends with Current Population Survey Data: A View From Inside the Census Bureau Vaults'. *Review of Income and Wealth* **55**(1), 166–185.
- Burkhauser, R. V., S. Feng, S. P. Jenkins, and J. Larrimore: 2012, 'Recent Trends in Top Income Shares in the United States: Reconciling Estimates from March CPS and IRS Tax Return Data'. *Review of Economics and Statistics* **94**(2), 371–388.
- Canberra Group: 2011, *Handbook on Household Income Statistics*. United Nations.
- Card, D.: 1999, 'The causal effect of education on earnings'. In: O. C. Ashenfelter and D. Card (eds.): *Handbook of Labor Economics*, Vol. 3A. Elsevier, pp. 1801–1863.
- Card, D., J. Heining, and P. Kline: 2013, 'Workplace Heterogeneity and the Rise of West German Wage Inequality'. *The Quarterly Journal of Economics* **128**(3), 967–1015.
- Carr-Hill, R.: 2013, 'Missing Millions and Measuring Development Progress'. *World Development* **46**(0), 30–44.

- Caselli, F. and W. J. Coleman: 2006, 'The World Technology Frontier'. *The American Economic Review* **96**(3), 499–522.
- CBO: 2012, 'The Distribution of Household Income and Federal Taxes, 2008 and 2009'. Congressional Budget Office Report 4441.
- Chen, S. and M. Ravallion: 2004, 'How Have the World's Poorest Fared since the Early 1980s?'. *The World Bank Research Observer* **19**(2), 141–169.
- Chen, S. and M. Ravallion: 2010a, 'China Is Poorer Than We Thought, But No Less Successful in the Fight against Poverty'. In: S. Anand, P. Segal, and J. Stiglitz (eds.): *Debates on the Measurement of Global Poverty*. Oxford University Press.
- Chen, S. and M. Ravallion: 2010b, 'The Developing World is Poorer than We Thought, But No Less Successful in the Fight Against Poverty'. *The Quarterly Journal of Economics* **125**(4), 1577–1625.
- Chetty, R. and E. Saez: 2005, 'Dividend Taxes and Corporate Behavior: Evidence from the 2003 Dividend Tax Cut'. *The Quarterly Journal of Economics* **120**(3), 791–833.
- Chotikapanich, D., W. E. Griffiths, D. S. P. Rao, and V. Valencia: 2012, 'Global Income Distributions and Inequality, 1993 and 2000: Incorporating Country-Level Inequality Modeled with Beta Distributions'. *The Review of Economics and Statistics* **94**(1), 52–73.
- Chotikapanich, D., R. Valenzuela, and D. Prasada Rao: 1997, 'Global and regional inequality in the distribution of income: Estimation with limited and incomplete data'. *Empirical Economics* **22**(4), 533–546.
- Christian, C. W.: 1994, 'Voluntary Compliance with the Individual Income Tax: Results from the 1988 TCMP Study'. *The IRS Research Bulletin* **1500**, 35. United States Department of the Treasury Internal Revenue Service.
- Clark, C.: 1951, *The Conditions of Economic Progress*. McMillan, 2nd edition.
- Cowell, F. A.: 2009, *Measuring Inequality*. Oxford University Press.
- Cowell, F. A. and E. Flachaire: 2007, 'Income distribution and inequality measurement: The problem of extreme values'. *Journal of Econometrics* **141**(2), 1044–1072.
- Credit Suisse: 2013, *Global Wealth Report Databook*.
- Criscuolo, C.: 2000, 'Employer Size - Wage Effect: A Critical Review and an Econometric Analysis'. Working paper, University of Siena.
- Curran, L. B., H. Wolman, E. W. Hill, and K. Furdell: 2008, 'Poverty, Programs, and Prices: How Adjusting for Costs of Living Would Affect Federal Benefit Eligibility'. Working paper, Brookings Institution Metropolitan Policy Program.
- Dardanoni, V. and P. Lambert: 2001, 'Horizontal inequity comparisons'. *Social Choice and Welfare* **18**(4), 799–816.
- Datt, G. and M. Ravallion: 1992, 'Growth and redistribution components of changes in poverty measures: A decomposition with applications to Brazil and India in the 1980s'. *Journal of Development Economics* **38**(2), 275–295.
- Davies, J. B., S. Sandström, A. Shorrocks, and E. N. Wolff: 2011, 'The Level and Distribution of Global Household Wealth'. *The Economic Journal* **121**(551), 223–254.

- Deaton, A.: 2001, 'Counting the World's Poor: Problems and Possible Solutions'. *The World Bank Research Observer* **16**(2), 125–147.
- Deaton, A.: 2005, 'Measuring Poverty in a Growing World (Or Measuring Growth in a Poor World)'. *The Review of Economics and Statistics* **87**(1), 1–19.
- Deaton, A.: 2010, 'Price Indexes, Inequality, and the Measurement of World Poverty'. *American Economic Review* **100**(1), 5–34.
- Deaton, A. and B. Aten: 2014, 'Trying to understand the PPPs in ICP 2011: why are the results so different?'. Working paper, Princeton University.
- Deaton, A. and O. Dupriez: 2011a, 'Purchasing Power Parity Exchange Rates for the Global Poor'. *American Economic Journal: Applied Economics* **3**(2), 137–66.
- Deaton, A. and O. Dupriez: 2011b, 'Spatial price differences within large countries'. Working paper, Princeton University.
- Deaton, A. and A. Heston: 2010, 'Understanding PPPs and PPP-based National Accounts'. *American Economic Journal: Macroeconomics* **2**(4), 1–35.
- Decancq, K.: 2009, 'Copula-based Measurement of Dependence Between Dimensions of Well-being'. Working paper, University of York. HEDG.
- Deckers, T., A. Falk, and H. Schildberg-Hörisch: 2013, 'Nominal or Real? The Impact of Regional Price Levels on Satisfaction with Life'. IZA Discussion Paper 7345.
- Deininger, K. and L. Squire: 1996, 'A New Data Set Measuring Income Inequality'. *World Bank Economic Review* **10**(3), 565–591.
- Del Río, C. and J. Ruiz-Castillo: 2001, 'Intermediate Inequality and Welfare: The Case of Spain, 1980-81 to 1990-91'. *Review of Income and Wealth* **47**(2), 221–37.
- DESTATIS: 2013, *Turnusmäßige Überarbeitung des Verbraucherpreisindex 2013*. Statistisches Bundesamt/Federal Statistics Office.
- Devarajan, S.: 2013, 'Africa's Statistical Tragedy'. *Review of Income and Wealth* **59**, 9–15.
- Diamond, R.: 2013, 'The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000'. Working paper, Stanford University.
- Dowrick, S. and M. Akmal: 2005, 'Contradictory Trends in Global Income Inequality: A Tale of Two Biases'. *Review of Income and Wealth* **51**(2), 201–229.
- Dreger, C. and R. Kosfeld: 2010, 'Do Regional Price Levels Converge?'. *Jahrbücher für Nationalökonomie und Statistik* **230**(3), 274–286.
- Dustmann, C., J. Ludsteck, and U. Schönberg: 2009, 'Revisiting The German Wage Structure'. *Quarterly Journal of Economics* **124**(2), 843–881.
- Dynan, K. E., J. Skinner, and S. P. Zeldes: 2004, 'Do the Rich Save More?'. *Journal of Political Economy* **112**(2), 397–444.
- Edsall, T. B.: 2013, 'What if We're Looking at Inequality the Wrong Way?'. *New York Times* (26 June 2013; accessed 21 February 2014).

- Elbers, C., P. Lanjouw, J. Mistiaen, B. Ozler, and K. Simler: 2005, 'Are Neighbors Equal? Estimating Local Inequality in Three Developing Countries,'. In: R. Kanbur and A. J. Venables (eds.): *Spatial Inequality and Development*. Oxford University Press.
- Elsby, M. W. L., B. Hobijn, and A. Sahin: 2013, 'The Decline of the U.S. Labor Share'. *Fall 2013 Brookings Panel on Economic Activity*.
- Eltető, O. and P. Köves: 1964, 'On a Problem of Index Number Computation Relating to International Comparison.' *Statisztikai Szemle* **42**, 507–518.
- Fajnzylber, P., W. Maloney, and G. M. Rojas: 2006, 'Microenterprise Dynamics in Developing Countries: How Similar are They to Those in the Industrialized World? Evidence from Mexico'. *The World Bank Economic Review* **20**(3), 389–419.
- Falco, P.: 2014, 'Does risk matter for occupational choices? Experimental evidence from an African labour market'. *Labour Economics* **28**, 96–109.
- Falco, P., A. Kerr, N. Rankin, J. Sandefur, and F. Teal: 2011, 'The returns to formality and informality in urban Africa'. *Labour Economics* **18**(1), 23–31.
- Feenberg, D. R. and J. M. Poterba: 1993, 'Income Inequality and the Incomes of Very High-Income Taxpayers: Evidence from Tax Returns'. *Tax Policy and the Economy* **7**, 145–177.
- Feenberg, D. R. and J. M. Poterba: 2000, 'The Income and Tax Share of Very High-Income Households, 1960-1995'. *The American Economic Review* **90**(2), 264–270.
- Feinstein, C. H.: 1968, 'Changes in the Distribution of National Income in the United Kingdom Since 1860'. In: J. Marchal and B. Ducros (eds.): *The Distribution of National Income*. Macmillan.
- Feldman, N. E. and J. Slemrod: 2007, 'Estimating Tax Noncompliance with Evidence from Unaudited Tax Returns'. *The Economic Journal* **117**(518), 327–352.
- Feldstein, M. S.: 2008, 'Did Wages Reflect Growth in Productivity?'. Working Paper 13953, National Bureau of Economic Research.
- Fields, G. S.: 1975, 'Rural-urban migration, urban unemployment and underemployment, and job-search activity in LDCs'. *Journal of Development Economics* **2**, 165–187.
- Fisher, J. D., D. S. Johnson, and T. M. Smeeding: 2013, 'Measuring the Trends in Inequality of Individuals and Families: Income and Consumption'. *American Economic Review: Papers and Proceedings* **103**(3), 184–188.
- Frankel, D. M. and E. D. Gould: 2001, 'The Retail Price of Inequality'. *Journal of Urban Economics* **49**(2), 219–239.
- Frick, J. R.: 2008, 'SOEPmonitor 1984-2007: Time Series on selected Indicators about Living in Germany'. DIW Berlin.
- Förster, M. F., D. Jesuit, and T. Smeeding: 2005, 'Regional Poverty and Income Inequality in Central and Eastern Europe: Evidence from the Luxembourg Income Study'. In: R. Kanbur and A. J. Venables (eds.): *Spatial Inequality and Development*. Oxford University Press.
- Fuchs-Schündeln, N. and R. Izem: 2012, 'Explaining the low labor productivity in East Germany-A spatial analysis'. *Journal of Comparative Economics* **40**(1), 1–21.

- Fullerton, D.: 1996, ‘Comment on Joel Slemrod: High-Income Families and the Tax Changes of the 1980s’. In: M. S. Feldstein and J. M. Poterba (eds.): *Empirical Foundations of Household Taxation*. University of Chicago Press.
- Funkhouser, E.: 1998, ‘The importance of firm wage differentials in explaining hourly earnings variation in the large-scale sector of Guatemala’. *Journal of Development Economics* **55**(1), 115–131.
- Gartner, H.: 2005, ‘The imputation of wages above the contribution limit with the German IAB employment sample’. FDZ Methodenreport 2.
- GDB: 2006, ‘General Description Booklet for the 2006 Public Use Tax File’. <http://users.nber.org/taxsim/gdb/>.
- Gerschenkron, A.: 1947, ‘The Soviet Indices of Industrial Production’. *The Review of Economics and Statistics* **29**(4), 217–226.
- Glaeser, E. L. and J. D. Gottlieb: 2008, ‘The Economics of Place-Making Policies’. Conference draft, Brookings Papers on Economic Activity.
- Goebel, J., J. R. Frick, and M. M. Grabka: 2009, ‘Preisunterschiede mildern Einkommensgefälle zwischen West und Ost’. *DIW Wochenbericht* **51**.
- Goldberg, P. K. and N. Pavcnik: 2007, ‘Distributional Effects of Globalization in Developing Countries’. *Journal of Economic Literature* **45**(1), 39–82.
- Goldin, C. and L. F. Katz: 2007, ‘The Race between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005’. NBER Working Paper 12984.
- Gollin, D.: 2002, ‘Getting Income Shares Right’. *Journal of Political Economy* **110**(2), 458–474.
- Goolsbee, A.: 2000, ‘What Happens When You Tax the Rich? Evidence from Executive Compensation’. *Journal of Political Economy* **108**(2), 352–378.
- Gordon, R. J. and I. Dew-Becker: 2008, ‘Controversies About The Rise of American Inequality: A Survey’. *NBER Working Paper Series* **13982**.
- Gottschalk, P. and S. Danziger: 2005, ‘Inequality of Wage Rates, Earnings and Family Income in the United States, 1975-2002’. *Review of Income and Wealth* **51**(2), 231–254.
- Goux, D. and E. Maurin: 1999, ‘Persistence of Interindustry Wage Differentials: A Reexamination Using Matched Worker-Firm Panel Data’. *Journal of Labor Economics* **17**(3), 492–533.
- Grimm, M.: 2007, ‘Removing the anonymity axiom in assessing pro-poor growth’. *The Journal of Economic Inequality* **5**(2), 179–197.
- Grün, C. and S. Klasen: 2008, ‘Growth, inequality, and welfare: comparisons across space and time’. *Oxford Economic Papers* **60**(2), 212–236.
- Groves, R. M. and M. Couper: 1998, *Nonresponse in Household Interview Surveys*. Wiley.
- Guerrieri, V., D. Hartley, and E. Hurst: 2013, ‘Endogenous gentrification and housing price dynamics’. *Journal of Public Economics* **100**, 45–60.

- Haig, R. M.: 1921, 'The Concept of Income - Economic and Legal Aspects'. In: R. M. Haig (ed.): *The Federal Income Tax*. Columbia University Press.
- Handbury, J.: 2013, 'Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities'. Working paper, The Wharton School, University of Pennsylvania.
- Harris, J. R. and M. P. Todaro: 1970, 'Migration, Unemployment and Development: A Two-Sector Analysis'. *The American Economic Review* **60**(1), 126–142.
- Hart, P. E.: 1976, 'The Comparative Statics and Dynamics of Income Distributions'. *Journal of the Royal Statistical Society. Series A (General)* **139**(1), 108–125.
- Haskel, J., R. Z. Lawrence, E. E. Leamer, and M. J. Slaughter: 2012, 'Globalization and U.S. Wages: Modifying Classic Theory to Explain Recent Facts'. *Journal of Economic Perspectives* **26**(2), 119–40.
- Hausman, J. A. and J. M. Poterba: 1987, 'Household Behavior and the Tax Reform Act of 1986'. *Journal of Economic Perspectives* **1**(1), 101–119.
- Hayashi, F.: 2000, *Econometrics*. Princeton University Press, Princeton.
- Hill, R. J.: 1997, 'A Taxonomy of Multilateral Methods for Making International Comparisons of Prices and Quantities'. *Review of Income and Wealth* **43**(1), 49–69.
- Hill, R. J. and T. P. Hill: 2003, 'Expectations, Capital Gains, and Income'. *Economic Inquiry* **41**(4), 607–619.
- Hillringhaus, T. and A. Peichl: 2010, 'Die Messung von Armut unter Berücksichtigung regional divergierender Lebenshaltungskosten und öffentlicher Leistungen'. IZA Discussion Paper 5344.
- Hlasny, V. and P. Verme: 2013, 'Top incomes and the measurement of inequality in Egypt'. Policy Research Working Paper Series 6557, The World Bank.
- Holt, S.: 2006, 'The Earned Income Tax Credit at Age 30: What We Know'. Research brief, Brookings Institution.
- Hungerford, T. L.: 2010, 'The Redistributive Effect of Selected Federal Transfer and Tax Provisions'. *Public Finance Review* **38**(4), 450–472.
- Hungerford, T. L.: 2013, 'Changes in Income Inequality Among U.S. Tax Filers between 1991 and 2006: The Role of Wages, Capital Income, and Taxes'. Working paper, Congressional Research Service.
- IRS: 2010, 'General Description Booklet for the 2006 Public Use Tax File'. *Statistics of Income Division, Internal Revenue Service*.
- Jäntti, M., M. Riihelä, R. Sullström, and M. Tuomala: 2010, 'Trends in Top Income Shares in Finland'. In: A. B. Atkinson and T. Piketty (eds.): *Top Incomes: A Global Perspective*. Oxford University Press.
- Jäntti, M., E. Sierminska, and P. V. Kerm: 2012, 'Modelling the joint distribution of income and wealth'. Presented at 32nd general conference, International Association for Research in Income and Wealth (IARIW).

- Johnson, D. G.: 1954, 'The Functional Distribution of Income in the United States, 1850-1952'. *The Review of Economics and Statistics* **36**(2), 175–182.
- Johnson, R., J. Nunns, J. Rohaly, E. Toder, and R. Williams: 2011, 'Why Some Tax Units Pay No Income Tax'. Working paper, Urban-Brookings Tax Policy Center.
- Johnston, R., M. McKinney, and T. Stark: 1996, 'Regional Price Level Variations and Real Household Incomes in the United Kingdom, 1979/80-1993'. *Regional Studies* **30**(6), 567–578.
- Jolliffe, D.: 2006, 'Poverty, Prices, and Place: How Sensitive is the Spatial Distribution of Poverty to Cost of Living Adjustments?'. *Economic Inquiry* **44**(2), 296–310.
- Kaplan, S. N. and J. Rauh: 2010, 'Wall Street and Main Street: What Contributes to the Rise in the Highest Incomes?'. *Review of Financial Studies* **23**(3), 1004–1050.
- Khamis, S. H.: 1972, 'A New System of Index Numbers for National and International Purposes'. *Journal of the Royal Statistical Society. Series A (General)* **135**(1), 96–121.
- Kingdon, G., J. Sandefur, and F. Teal: 2006, 'Labour Market Flexibility, Wages and Incomes in Sub-Saharan Africa in the 1990s'. *African Development Review* **18**(3), 392–427.
- Knight, J. and S. Ding: 2012, *China's Remarkable Economic Growth*. Oxford University Press.
- Kolenikov, S. and A. Shorrocks: 2005, 'A Decomposition Analysis of Regional Poverty in Russia'. *Review of Development Economics* **9**(1), 25–46.
- Kolm, S.-C.: 1976, 'Unequal Inequalities I'. *Journal of Economic Theory* **12**(3), 416–42.
- Kopczuk, W. and E. Saez: 2004, 'Top Wealth Shares in the United States, 1916-2000: Evidence from Estate Tax Returns'. *National Tax Journal* **57**(2), 445–487.
- Kopczuk, W., E. Saez, and J. Song: 2010, 'Earnings Inequality and Mobility in the United States: Evidence from Social Security Data since 1937'. *Quarterly Journal of Economics* **125**(1), 91–128.
- Korinek, A., J. A. Mistiaen, and M. Ravallion: 2006, 'Survey nonresponse and the distribution of income'. *The Journal of Economic Inequality* **4**(1), 33–55.
- Kosfeld, R., H.-F. Eckey, and M. Schüßler: 2009, 'Ökonometrische Messung regionaler Preisniveaus auf der Basis örtlich beschränkter Erhebungen'. *RatSWD Research Notes* **33**, 1–31.
- Krause, P.: 1994, 'Armut im Wohlstand: Betroffenheit und Folgen'. *DIW Diskussionspapier* **88**.
- Krueger, D. and F. Perri: 2006, 'Does Income Inequality Lead to Consumption Inequality? Evidence and Theory'. *The Review of Economic Studies* **73**(1), 163–193.
- Kuznets, S.: 1955, 'Economic Growth and Income Inequality'. *The American Economic Review* **45**(1), 1–28.
- Leigh, A.: 2007, 'How Closely Do Top Income Shares Track Other Measures of Inequality?'. *The Economic Journal* **117**(524), 619–633.

- Leigh, A. and P. van der Eng: 2009, 'Inequality in Indonesia: What can we learn from top incomes?'. *Journal of Public Economics* **93**, 209–212.
- Lerman, R. I. and S. Yitzhaki: 1984, 'A note on the calculation and interpretation of the Gini index'. *Economics Letters* **15**, 363–368.
- Lerman, R. I. and S. Yitzhaki: 1985, 'Income Inequality Effects by Income Source: A New Approach and Applications to the United States'. *The Review of Economics and Statistics* **67**(1), 151–156.
- Lessmann, C.: 2014, 'Spatial inequality and development - Is there an inverted-U relationship?'. *Journal of Development Economics* **106**(0), 35–51.
- Li, H., L. Squire, and H.-f. Zou: 1998, 'Explaining International and Intertemporal Variations in Income Inequality'. *The Economic Journal* **108**(446), 26–43.
- Lofquist, D., T. Lugaila, M. O'Connell, and S. Feliz: 2012, 'Households and Families: 2010'. Census briefs C2010BR-14, US Census Bureau. <http://www.census.gov/prod/cen2010/briefs/c2010br-14.pdf>.
- Lucas, R. E. J.: 1978, 'On the Size Distribution of Business Firms'. *The Bell Journal of Economics* **9**(2), 508–523.
- Lydall, H.: 1968, *The Structure of Earnings*. Clarendon Press.
- MacKay, N.: 2011, 'London house prices are power-law distributed'. Working paper, University of York.
- Maddala, G. S.: 1983, *Limited-dependent and qualitative variables in econometrics*. Cambridge University Press, Cambridge.
- Maddison, A.: 2008, 'The West and the Rest in the World Economy: 1000-2030'. *World Economics* **9**(4), 75–99.
- Maloney, W. F.: 2004, 'Informality Revisited'. *World Development* **32**(7), 1159–1178.
- Mann, E. S. and J. J. Salvo: 1984, 'Linking IRS and Census Data to Produce Intercensal Income Estimates: An Exploration'. In: *Statistics of Income and Related Administrative Record Research: 1984*. Internal Revenue Service.
- Meade, J. E.: 1964, *Efficiency, equality and the ownership of property*. Allen & Unwin.
- Melitz, M. J.: 2003, 'The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity'. *Econometrica* **71**(6), 1695–1725.
- Milanovic, B.: 2002, 'True World Income Distribution, 1988 and 1993: First Calculation Based on Household Surveys Alone'. *The Economic Journal* **112**(476), 51–92.
- Milanovic, B.: 2005, *Worlds Apart: Measuring International and Global Inequality*. Princeton University Press.
- Milanovic, B.: 2012, 'Global inequality recalculated and updated: the effect of new PPP estimates on global inequality and 2005 estimates'. *The Journal of Economic Inequality* **10**, 1–18.
- Milanovic, B. and S. Yitzhaki: 2002, 'Decomposing World Income Distribution: Does The World Have A Middle Class?'. *Review of Income and Wealth* **48**(2), 155–178.

- Mincer, J.: 1974, *Schooling, Experience and Earnings*. Columbia University Press, New York.
- Minoiu, C. and S. Reddy: 2012, ‘Kernel density estimation on grouped data: the case of poverty assessment’. *The Journal of Economic Inequality* pp. 1–27.
- Mistiaen, J. A. and M. Ravallion: 2003, ‘Survey compliance and the distribution of income’. Policy Research Working Paper Series 2956, The World Bank.
- Mogstad, M., A. Langørgen, and R. Aaberge: 2007, ‘Region-specific versus country-specific poverty lines in analysis of poverty’. *The Journal of Economic Inequality* **5**(1), 115–122.
- Moretti, E.: 2013, ‘Real Wage Inequality’. *American Economic Journal: Applied Economics* **5**(1), 65–103.
- Morival, E.: 2011, ‘Top incomes and racial inequality in South Africa: Evidence from tax statistics and household surveys 1993-2008’. Master’s thesis, Paris School of Economics.
- Mortensen, D.: 2003, *Wage dispersion: why are similar workers paid differently*. MIT Press, Cambridge.
- Moulton, B. R.: 1995, ‘Interarea indexes of the cost of shelter using hedonic quality adjustment techniques’. *Journal of Econometrics* **68**(1), 181–204.
- Muellbauer, J.: 1974, ‘Inequality Measures, Prices and Household Composition’. *The Review of Economic Studies* **41**(4), 493–504.
- Murphy, K. J.: 2002, ‘Explaining Executive Compensation: Managerial Power versus the Perceived Cost of Stock Options’. *The University of Chicago Law Review* **69**(3), 847–869.
- Nagel, T.: 2005, ‘The Problem of Global Justice’. *Philosophy and Public Affairs* **33**(2), 113–147.
- Nelsen, R. B.: 2006, *An Introduction to Copulas*. Springer.
- Nuxoll, D. A.: 1994, ‘Differences in Relative Prices and International Differences in Growth Rates’. *The American Economic Review* **84**(5), 1423–1436.
- OECD: 2008, ‘Growing Unequal? Income Distribution and Poverty in OECD Countries.’. Working paper, OECD Publishing.
- Ohnishi, T., T. Mizuno, C. Shimizu, and T. Watanabe: 2011, ‘On the Evolution of the House Price Distribution’. Working paper, Center on Japanese Economy and Business, Columbia University.
- Oi, W. Y. and T. L. Idson: 1999, ‘Chapter 33 Firm size and wages’. In: O. C. Ashenfelter and D. Card (eds.): *Handbook of Labor Economics*, Vol. 3, Part 2 of *Handbook of Labor Economics*. Elsevier, pp. 2165–2214.
- Piketty, T. and E. Saez: 2001, ‘Income Inequality in the United States, 1913-1998’. Working Paper 8467, National Bureau of Economic Research.
- Piketty, T. and E. Saez: 2003, ‘Income Inequality in the United States, 1913-1998’. *The Quarterly Journal of Economics* **118**(1), 1–39.
- Piketty, T. and E. Saez: 2007a, ‘How Progressive is the U.S. Federal Tax System? A Historical and International Perspective’. *Journal of Economic Perspectives* **21**(1), 3–24.

- Piketty, T. and E. Saez: 2007b, 'Income and Wage Inequality in the United States 1913-2002'. In: A. B. Atkinson and T. Piketty (eds.): *Top Incomes over the Twentieth Century: A Contrast Between Continental European and English-Speaking Countries*. Oxford University Press.
- Piketty, T., E. Saez, and S. Stantcheva: 2011, 'Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities'. NBER Working Paper 17616, NBER.
- Pinkovskiy, M. L.: 2013, 'World welfare is rising: Estimation using nonparametric bounds on welfare measures'. *Journal of Public Economics* **97**(0), 176–195.
- Pischke, J.-S. and T. von Wachter: 2008, 'Zero Returns to Compulsory Schooling in Germany: Evidence and Interpretation'. *Review of Economics and Statistics* **90**(3), 592–598.
- Pissarides, C. A. and G. Weber: 1989, 'An expenditure-based estimate of Britain's black economy'. *Journal of Public Economics* **39**(1), 17–32.
- Pittau, M. G., R. Zelli, and R. Massari: 2011, 'Do Spatial Price Indices Reshuffle the Italian Income Distribution?'. *Modern Economy* **2**, 259–265.
- Pogge, T. W.: 2002, *World Poverty and Human Rights: Cosmopolitan Responsibilities and Reforms*. Polity Press.
- Quinn, C.: 2007, 'Using copulas to measure association between ordinal measures of health and income'. Working paper, University of York. HEDG.
- Rankin, N., J. Sandefur, and F. Teal: 2010, 'Learning and Earning in Africa: Where are the returns to education high?'. Working paper, University of Oxford.
- Rao, V.: 2000, 'Price Heterogeneity and "Real" Inequality: A Case Study of Prices and Poverty in Rural South India'. *Review of Income and Wealth* **46**(2), 201–211.
- Ravallion, M.: 2004, 'Competing Concepts of Inequality in the Globalization Debate'. In: S. M. Collins and C. Graham (eds.): *Brookings Trade Forum 2004: Globalization, Poverty, and Inequality*. Brookings Institution Press.
- Ravallion, M.: 2008, 'A Global Perspective on Poverty in India'. *Economic and Political Weekly* **43**(43), 33–37.
- Ravallion, M.: 2010a, 'A reply to Reddy and Pogge'. In: S. Anand, P. Segal, and J. E. Stiglitz (eds.): *Debates on the Measurement of Global Poverty*. Oxford University Press.
- Ravallion, M.: 2010b, 'Understanding PPPs and PPP-based National Accounts: Comment'. *American Economic Journal: Macroeconomics* **2**(4), 46–52.
- Ravallion, M.: 2014a, 'An Exploration of the International Comparison Program's New Global Economic Landscape'. Working paper, Georgetown University.
- Ravallion, M.: 2014b, 'The Luxembourg Income Study'. Working paper, ECINEQ.
- Ravallion, M. and S. Chen: 2003, 'Measuring pro-poor growth'. *Economics Letters* **78**(1), 93–99.
- Ravallion, M. and S. Chen: 2007, 'China's (uneven) progress against poverty'. *Journal of Development Economics* **82**(1), 1–42.

- Ravallion, M., S. Chen, and P. Sangraula: 2007, 'New Evidence on the Urbanization of Global Poverty'. *Population and Development Review* **33**(4), 667–701.
- Reddy, S. G. and T. Pogge: 2010, 'How Not to Count the Poor'. In: S. Anand, P. Segal, and J. E. Stiglitz (eds.): *Debates on the Measurement of Global Poverty*. Oxford University Press.
- Rienzo, C.: 2010, 'Real Wages, Wage Inequality and the Regional Cost-of-living in the UK'. Working paper, Royal Holloway.
- Ringuedé, S.: 1998, 'An efficiency wage model for small firms: firm size and wages'. *Economics Letters* **59**, 263–268.
- Riphahn, R. T. and D. D. Schnitzlein: 2011, 'Wage Mobility in East and West Germany'. IZA Discussion Paper 6246.
- Roback, J.: 1982, 'Wages, Rents, and the Quality of Life'. *Journal of Political Economy* **90**(6), 1257–1278.
- Roine, J. and D. Waldenström: 2008, 'The evolution of top incomes in an egalitarian society: Sweden, 1903-2004'. *Journal of Public Economics* **92**, 366–387.
- Roine, J. and D. Waldenström: 2012, 'On the Role of Capital Gains in Swedish Income Inequality'. *Review of Income and Wealth* **58**(3), 569–587.
- Roodman, D.: 2009a, 'How to do xtabond2: An introduction to difference and system GMM in Stata'. *The Stata Journal* **9**(1), 86–136.
- Roodman, D.: 2009b, 'Practitioners' corner: A note on the theme of too many instruments'. *Oxford Bulletin of Economics and Statistics* **71**(1), 135–158.
- Roos, M.: 2006, 'Earnings Disparities in Unified Germany: Nominal versus Real'. *Jahrbuch für Regionalwissenschaft* **26**, 171–189.
- Rosen, H. S. and T. Gayer: 2010, *Public finance*. McGraw-Hill, 8th edition.
- Ruiz-Castillo, J.: 1998, 'A Simplified Model for Social Welfare Analysis: An Application to Spain, 1973-74 to 1980-81'. *Review of Income and Wealth* **44**(1), 123–141.
- Saez, E.: 2004, 'Reported Incomes and Marginal Tax Rates, 1960-2000: Evidence and Policy Implications'. *Tax Policy and the Economy* **18**, 117–173.
- Saez, E.: 2013, 'Income inequality in the United States, 1913-2012'. Updated Tables and Figures on <http://elsa.berkeley.edu/saez/>, accessed: 07.09.2014.
- Saez, E., J. Slemrod, and S. H. Giertz: 2012, 'The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review'. *Journal of Economic Literature* **50**(1), 3–50.
- Saez, E. and M. R. Veall: 2007, 'The Evolution of High Incomes in Canada, 1920-2000'. In: A. B. Atkinson and T. Piketty (eds.): *Top Incomes over the Twentieth Century. A Contrast Between Continental European and English-Speaking Countries*. Oxford University Press.
- Sala-i Martin, X.: 2006, 'The World Distribution of Income: Falling Poverty and ... Convergence, Period'. *The Quarterly Journal of Economics* **121**(2), 351–397.
- Samuelson, P. A.: 1994, 'Facets of Balassa-Samuelson Thirty Years Later'. *Review of International Economics* **2**(3), 201–226.

- Sandefur, J., P. Serneels, and F. Teal: 2007, 'Poverty and Earnings Mobility in Three African Countries'. In: P. Paci and P. Serneels (eds.): *Employment and Shared Growth : Rethinking the Role of Labor Mobility for Development*. World Bank.
- Schmieder, J. F., T. von Wachter, and S. Bender: 2012, 'The Effects of Extended Unemployment Insurance Over the Business Cycle: Evidence from Regression Discontinuity Estimates Over 20 Years'. *The Quarterly Journal of Economics* **127**(2), 701–752.
- Söderbom, M. and F. Teal: 2004, 'Size and efficiency in African manufacturing firms: evidence from firm-level panel data'. *Journal of Development Economics* **73**(1), 369–394.
- Söderbom, M., F. Teal, and A. Wambugu: 2005, 'Unobserved heterogeneity and the relation between earnings and firm size: evidence from two developing countries'. *Economics Letters* **87**(2), 153–159.
- Söderbom, M., F. Teal, A. Wambugu, and G. Kahyarara: 2006, 'The Dynamics of Returns to Education in Kenyan and Tanzanian Manufacturing'. *Oxford Bulletin of Economics and Statistics* **68**(3), 261–288.
- Segal, P.: 2011, 'Resource Rents, Redistribution, and Halving Global Poverty: The Resource Dividend'. *World Development* **39**, 475–489.
- Shimer, R. and L. Smith: 2000, 'Assortative Matching and Search'. *Econometrica* **68**, 371–398.
- Shorrocks, A. and G. Wan: 2008, 'Ungrouping Income Distributions'. Working paper, UNU-WIDER.
- Shorrocks, A. F.: 1982, 'Inequality Decomposition by Factor Components'. *Econometrica* **50**(1), 193–211.
- SIAB: 2010, 'Sample of Integrated Labour Market Biographies 1975-2008'. *Research Data Centre of the German Federal Employment Agency at the Institute for Employment Research, Nuremberg*.
- Simons, H. C.: 1938, *Personal Income Taxation: the Definition of Income as a Problem of Fiscal Policy*. University of Chicago Press.
- Singer, P.: 2002, *One World: The Ethics of Globalization*. Yale University Press.
- Slemrod, J.: 1992, 'Do Taxes Matter? Lessons from the 1980's'. *American Economic Review Papers and Proceedings* **82**(2), 250–256.
- Slemrod, J.: 1996, 'High-Income Families and the Tax Changes of the 1980s: The Anatomy of Behavioral Response'. In: M. S. Feldstein and J. M. Poterba (eds.): *Empirical Foundations of Household Taxation*. University of Chicago Press.
- Slemrod, J.: 2007, 'Cheating Ourselves: The Economics of Tax Evasion'. *Journal of Economic Perspectives* **21**(1), 25–48.
- Slesnick, D. T.: 2002, 'Prices and Regional Variation in Welfare'. *Journal of Urban Economics* **51**(3), 446–468.
- Smeeding, T. M. and J. P. Thompson: 2011, *Who Loses in the Downturn? Economic Crisis, Employment and Income Distribution (Research in Labor Economics, Volume 32)*, Chapt. Recent Trends in Income Inequality, pp. 1–50. Emerald Group Publishing Limited.

- SOEP Group: 2014, 'SOEP 2012 -SOEPmonitor Person 1984-2012 (SOEP v29)'. *SOEP Survey Papers* **193**.
- Solow, R.: 2014, 'Correspondence: The One Percent: Robert Solow, N. Gregory Mankiw'. *Journal of Economic Perspectives* **28**(1), 243–248.
- Somekh, B.: 2012, 'The Effect of Income Inequality on Price Dispersion'. Working paper, University of Haifa.
- Srinivasan, T. N.: 2001, 'Comment on "Counting the World's Poor," by Angus Deaton'. *The World Bank Research Observer* **16**(2), 157–168.
- Stiglitz, J. E., A. Sen, and J.-P. Fitoussi: 2009, 'Report by the Commission on the Measurement of Economic Performance and Social Progress'. Final report, <http://www.stiglitz-sen-fitoussi.fr/>.
- STMWIVT: 2003, *Die reale Kaufkraft in Bayern 2002: Zwischenörtliche Preis- und Einkommensunterschiede*. Bayerisches Staatsministerium für Wirtschaft, Verkehr und Technologie.
- Ströhl, G.: 1994, 'Zwischenörtlicher Vergleich des Verbraucherpreisniveaus in 50 Städten'. *Wirtschaft und Statistik* **6**, 415.
- Strudler, M., T. Petska, and R. Petska: 2005, 'Further Analysis of the Distribution of Income and Taxes, 1979-2002'. Working paper, Internal Revenue Service.
- Sutton, J.: 1997, 'Gibrat's Legacy'. *Journal of Economic Literature* **35**(1), 40–59.
- Székely, M. and M. Hilgert: 1999, 'What's Behind the Inequality We Measure: An Investigation Using Latin American Data'. *Research Department Working Paper Inter-American Development Bank*.
- Szulc, B.: 1964, 'Indices for Multiregional Comparisons'. *Przegląd Statystyczny* **3**, 239–254.
- Teal, F.: 2000, 'Real wages and the demand for skilled and unskilled male labour in Ghana's manufacturing sector: 1991-1995'. *Journal of Development Economics* **61**(2), 447–461.
- Teal, F.: 2008, 'The Ghana and Tanzania Urban Household Panel Surveys: 2004-2006'. Working paper, Centre for the Study of African Economies, University of Oxford.
- Teal, F.: 2011, 'The price of labour and understanding the causes of poverty'. *Labour Economics* **18**(1), 7–15.
- Thompson, J. P. and T. M. Smeeding: 2013, 'Inequality and Poverty in the United States: the Aftermath of the Great Recession'. Finance and economics discussion series, Federal Reserve Board.
- Trivedi, P. K. and D. M. Zimmer: 2005, 'Copula Modeling: An Introduction for Practitioners'. *Foundations and Trends in Econometrics* **1**(1), 1–111.
- United Nations: 2013, 'The 0.7 percent target: An in-depth look'. <http://www.unmillenniumproject.org/press/07.htm>, Accessed 28.05.2013.
- van Kerm, P.: 2009, 'Income mobility profiles'. *Economics Letters* **102**(2), 93–95.
- van Kerm, P.: 2010, 'sgini: Generalized Gini and Concentration coefficients (with factor decomposition) in Stata'. CEPS/INSTEAD, Luxembourg.

- Velenchik, A. D.: 1997, 'Market power, firm performance and real wage growth in Zimbabwean manufacturing'. *World Development* **25**(5), 749–762.
- Wang, S., Z. Yang, and H. Liu: 2011, 'Impact of urban economic openness on real estate prices: Evidence from thirty-five cities in China'. *China Economic Review* **22**(1), 42–54.
- Warner, D., D. P. Rao, W. E. Griffiths, and D. Chotikapanich: 2012, 'Global inequality: Levels and Trends, 1993-2005. How sensitive are these to the choice of PPPs and Real Incomes measures?'. Presented at 32nd general conference, International Association for Research in Income and Wealth (IARIW).
- Weiss, A. and H. J. Landau: 1984, 'Wages, Hiring Standards, and Firm Size'. *Journal of Labor Economics* **2**(4), 477–499.
- Windmeijer, F.: 2005, 'A finite sample correction for the variance of linear efficient two-step GMM estimators'. *Journal of Econometrics* **126**(1), 25–51.
- Wingfield, D., D. Fenwick, and K. Smith: 2005, 'Relative Regional Consumer Price Levels in 2004'. *Economic Trends* **615**, 36–46. Office for National Statistics.
- Winglee, M., R. Valliant, J. Clark, Y. Lim, M. Weber, and M. Strudler: 2002, 'Assessing Disclosure Protection for a SOI Public Use File'. *Proceedings of the American Statistical Association*.
- Winter-Ebmer, R. and J. Zweimüller: 1999, 'Firm-Size Wage Differentials in Switzerland: Evidence from Job-Changers'. *The American Economic Review Papers and Proceedings* **89**(2), 89–93.
- Wolff, E. N. and A. Zacharias: 2009, 'Household wealth and the measurement of economic well-being in the United States'. *The Journal of Economic Inequality* **7**(2), 83–115.
- Wood, A.: 1995, 'How Trade Hurt Unskilled Workers'. *Journal of Economic Perspectives* **9**(3), 57–80.
- World Bank: 2012, *World Development Report 2013: Jobs*. Washington, DC: World Bank.
- Yagan, D.: 2013, 'Capital Tax Reform and the Real Economy: The Effects of the 2003 Dividend Tax Cut'. Working paper, Harvard University.
- Zucman, G.: 2013, 'The Missing Wealth of Nations: Are Europe and the U.S. net Debtors or net Creditors?'. *The Quarterly Journal of Economics*.