GEORG-AUGUST UNIVERSITÄT GÖTTINGEN

DOKTORARBEIT

Economic Inequality: Causes, Consequences, and Measurement Issues. An Empirical Contribution.

Dissertation zur Erlangung des wirtschaftswissenschaftlichen Doktorgrades der Wirtschaftswissenschaftlichen Fakultät der Universität Göttingen

vorgelegt von

Nathalie Scholl

aus

Ulm

Göttingen, 2016

Erstgutachter:	Prof. Stephan Klasen, Ph.D.
Zweitgutachter:	Prof. Dr. Axel Dreher
Drittprüferin:	Prof. Inmaculada Martínez-Zarzoso, Ph.D.
Tag der mündlichen Prüfung:	27. Juni 2016

Contents

List of Figures

List of Tables

Chapter 1

Introduction

> **"There should exist among the citizens neither extreme poverty nor again excessive wealth, for both are productive of great evil." (Plato)**

Inequality and fairness have always been at the heart of any political concept aiming at the well-being of the people. The idea of equality is as old as democracy itself, and when selecting an inspiring quotation to start this dissertation with, I had the choice between statements by political thinkers from Ancient Greece's Plato to the father of modern India's democracy, Mahatma Gandhi, to German philosopher Karl Marx. In the end, I picked Plato's words because they capture many aspects which are still highly relevant to the debate on inequality - used here in the sense of monetary inequality, comprising everything from wage to income to wealth - in economics today.

The mentioning of both extreme poverty and excessive wealth relates to how inequality can stem from different parts of the income distribution. Poverty and inequality are distinct concepts and fields of research, but poverty in practice always implies inequality, and many of the arguments brought forward against inequality pertain to the people at the bottom of the income distribution. Most of these arguments have to do with equality of opportunity and touch upon other concepts closely related to, but different from inequality, such as social mobility. More recently, there has been a renewed interest in the top of the income distribution with the publication of Thomas Piketty's widely received "Capital in the Twenty-First Century" in 2014, which has also reached the popular debate and sparked interest in the topic of inequality as a whole. While there is a general consensus that (absolute) poverty has decreased substantially over the last few decades, Piketty argues that inequality has been rising steadily since over the second half of the 21th century in the industrialized countries. This is due to developments at the top end of the wealth distribution, where incomes have become worrisomely high and continue to rise at an increasing rate if no counteracting measures are taken. While these arguments have not remained uncontested (see, e.g., critiques by Stiglitz[1] and numerous newspaper articles), he still makes the same point that Plato made some 2500 years earlier: that excessive inequality will have adverse consequences. More specifically, too-high inequality will lead to an imbalance of power between the rich and the rest of society, thereby eroding the

[1] http://opinionator.blogs.nytimes.com/category/the-great-divide/?_r=1, accessed on April 17, 2016

fundamentals of democracy. This may lead to social unrest, and ultimately also hurt the material well-being of society as a whole, by hampering economic growth.

There are also reasons to believe that inequality is not a bad thing and can indeed lead to an increase, rather than reduction, in overall wealth. Philosopher Thomas Steven Molnar, capturing the essence of the US-American conservative view stated that "passion for equality blinds the utopian to the fact that society, as a whole, is based on inequality of men in two respects: the inventor, the innovator, the exceptional man creates something new and insures continuous progress; the others emulate his work or merely improve their own lot by benefiting from his creativity. Now, to deny to this exceptional man the extra compensation [...] is to extirpate his inventiveness." (Molnar 1972: 153). The effort certain people put into their work is what leads to the creation of an ever-increasing amount of material goods and resources in the first place, which also entitles them to "have a bigger piece" of the same. Moreover, knowing that they will get a bigger piece as a result of their hard work is exactly what motivates them to exert more effort in the first place.

Which view is "more right" in terms of the effect on a society's overall and long-term material well-being is ultimately an empirical question, and one to which this thesis aims to contribute.

Something that is implicit in the above discussion is the focus on inequality within countries. That is, judgments about the extent of inequality rely on comparisons of individuals from the same country. This is not to say that global inequality does not matter. With the ever-increasing access to, and availability of information through the rapid progress in communication technology, people everywhere become more and more aware of the living conditions of individuals in other countries, and comparisons are no longer made only with people in their own country. There is also a broad consensus that inequality between countries is still much larger than inequality within countries, with estimations of the contribution of between-country inequality to overall global inequality ranging from 55% to 90% (Anand and Segal 2008). The single most important determinant of a person's material welfare continues to be his or her birthplace (World Development Report 2009). The question of why some countries are richer than others can, of course, also be related to within-country inequality in the long run. If inequality within countries has repercussions on economic growth, it will thereby also affect future inequality between countries. Nevertheless, so long as policies are implemented at the level of the nation state, it is more relevant to know what the causes and consequences of inequality are within countries.

What are the factors which cause inequality to rise or fall within a country over time, and what are the implications of this higher or lower inequality for a country's economic prosperity? While it would be presumptuous to seek for universal answers to these questions within the purview of these 150 or so pages, this dissertation tries to provide partial explanations by focusing on specific aspects. Restricting my attention to one of the arguably most relevant consequences of inequality in terms of its long-run impact on welfare, I delve into the still unresolved debate on whether inequality is conducive to economic growth. I then go deeper into the issue of measurement of inequality, concentrating on a measure of wage inequality for one of the main sectors of the economy, the manufacturing industry. The measurement debate is not a trivial one: both causes and consequences of inequality cannot be explored without going into the details of inequality measures, and the availability of appropriate data across countries and over time is a prerequisite for empirical analyses on the topic. In fact, one of the reasons why there is no consensus as to whether inequality between countries has been increasing or decreasing over the last 50 years is the variety of not only measures but also measurement concepts that are used in different studies. These diverse measures are sensitive to distinct parts of the income distribution and hence can change rank orderings of countries, or of changes in inequality over time (Anand and Segal 2008). Similarly, one of the reasons why the literature is still inconclusive about the impact of inequality on economic growth is the lack of comparability due to the use of different data sources, as well as deficiencies in the underlying data sources themselves. Finally, I empirically look into one of the potential factors causing inequality, namely trade, and restrict my attention to one element of inequality, namely the distribution of wage incomes.

What remains to be said is that this dissertation does not entail any normative judgment regarding inequality, and whether it is desirable or not for reasons other than the ones highlighted above and in the three essays that follow. While all of the arguments mentioned thus far are concerned with the question of whether inequality is instrumental in increasing the overall material well-being of societies, there are of course also intrinsic reasons to value equality. The existence of inequality aversion has become a generally accepted fact in economics (see, e.g., Engelmann and Strobel 2004). There is also strong evidence from observational as well as experimental data that people's inequality aversion only refers to outcomes perceived as unfair, that is, inequality which arises due to factors beyond individual control (e.g., Cojocaru 2011, Alesina and La Ferrara 2005). Neurobiological findings suggest that people react differently to inequalities perceived as unfair (Cappelen et al. 2014), and there is even evidence from an experiment with monkeys which finds that the animals refuse to cooperate with their human experimenter when they receive unequal rewards for performed tasks (Brosnan and De Waal 2003). The mere existence of unfair inequality is thereby associated with a direct loss in welfare.

The remainder of this thesis is structured as follows: essay one starts by empirically investigating into one of arguably most important consequences of income inequality, economic growth. Essay two continues the debate on inequality data and measurement already touched upon in essay one while focusing on a very specific measure of wage inequality. Finally, despite the fact that the debate on the effects of inequality on economic growth remains unresolved, it is important to know which factors cause, or reinforce to it (and which ones do not). Essay three analyses one of the more frequently cited reasons in the public debate for rising (wage) inequality, namely trade.

Essay 1 is joint work with Stephan Klasen and revisits the inequality-growth relationship, which, despite a dizzying amount of theoretical and empirical studies on the topic, remains unresolved. Using an enhanced panel data set with improved inequality data and special attention to the role of transition countries, we replicate some of the most important contributions. We base our analysis on the specification of Forbes (2000), but also address the functional form concerns raised by Banerjee and Duflo (2003). The essay arrives at three main findings: First, similar to Forbes, we find a significant positive association between inequality and subsequent economic growth in the full sample, but this is entirely driven by transition (post-Soviet) countries. Second, this positive relationship in transition countries is not robust to the inclusion of separate time effects. Lastly, it therefore appears that this association is not causal but rather driven by the particular dynamics of the transition. Our finding is consistent with the claim that the relationship between inequality and growth emerges due to the particular timing of inequality and growth dynamics in transition countries. In particular, the rise in inequality in the 1990s coincided with a sharp output collapse, leading us to find an association between the large increase in inequality in the early 1990s and a growth recovery in the late 1990s. In sum, once the transition country dynamics are accounted for, we find no robust, systematic relationship between inequality and subsequent growth, neither for levels nor for changes in inequality. These results hold for different lag structures as well as in the medium- rather than the short term, and the empirical patterns observed are robust to the use of different data sets on inequality.

Essay 2 builds in part on the measurement debate already touched upon in essay 1. It introduces a newly constructed Theil index of between-sectoral manufacturing wage inequality and empirically tests whether the measure can serve as a basis for more general statements about the evolution of broader concepts of inequality, as argued by the authors of the very similar University of Texas Inequality Project (UTIP) index. Building on prior work of the UTIP, several concerns regarding the treatment of the raw data as well as important questions of internal and external validity are addressed. The index is based on sector-level data from the UNIDO Industrial Statistics for manufacturing, and I provide a detailed account of how the unbalanced raw data were treated in response to the lack of documentation for the UTIP measure. The newly computed index has the advantage of being consistently measured across countries and years, which makes it a valuable resource for empirical studies focusing on changes in the manufacturing structure within countries over long periods of time, such as the analysis conducted in essay 3. However, its narrow scope also restricts the applicability of the index for other, broader uses. I argue that the latter point is one of the main drawbacks of the index and present evidence that the generalizability from between-sectoral manufacturing wage inequality to overall income inequality is severely limited. This applies not only to the extent to which the index allows conjectures about the overall level of income inequality in a society, there is reason to also question the "internal" ability of the index to accurately reflect developments in manufacturing wage inequality. I therefore do not recommend using the index as a basis for inference about the development of broader concepts of wage- or income equality.

Essay 3 empirically tests whether trade has caused changes in wage inequality in developing countries, drawing on the measure of wage inequality constructed in essay 2. It builds on the empirical observation that since the expansion of world trade in the 1980s, measures of inequality have risen not only in developed countries, but also throughout the developing world. This stylized fact is contrary to the predictions of classical trade theory that in countries with high endowments of unskilled labor, their wages should rise relative to those of skilled labor. The essay tests the effects of trade on wage inequality in a differentiated panel framework where countries are classified according to their relative human capital endowments, constituting also the relevant comparative advantage in trade. Employing a newly constructed measure of technological change, an important source of omitted variable bias, not yet addressed in the literature, is removed. With the inclusion of this measure, several effects otherwise attributed to trade disappear, underscoring the importance of controlling for technological change. Technology transfer as well as technological change is found to take place particularly in industries and trade flows classified as medium-technology intensive. Evidence is also found for pure "trade"-effects, supporting the Heckscher-Ohlin predictions of the effects of trade on wage inequality once the heterogeneity of the trading partners and the traded goods is taken into account.

Chapter 2

Re-estimating the Relationship between Inequality and Growth

2.1 Introduction

The possible trade-off between inequality and growth has been investigated theoretically and empirically for decades. In the mid-1990s, the empirical debate was significantly enhanced by the availability of a much broader set of data on inequality across the world. Initially, the workhorse dataset was created by Deininger and Squire (DS1996) and used in a study by Deininger and Squire (1998) to show that, in a cross-section of countries, initial inequality (particularly of assets but, in some specifications, also of income) was associated with subsequent lower growth.

Ensuing debates focused on the one hand on weaknesses in the data, where Atkinson and Brandolini (2001) showed that the comparability and consistency of the (DS1996) data set was open to question. Since then, the World Income and Inequality Database (WIID) was created which significantly enhanced not only the coverage but also the transparency of the inequality data used. Many studies on inequality have since relied on this dataset, where some authors used regression-based adjustment methods to address inconsistency issues (e.g., Gruen and Klasen 2008, 2012; Easterly 2007). Nevertheless, the dataset remains heterogeneous in terms of the underlying monetary concept (covering not only different types of income (net, gross, wage incomes) but also consumption and expenditure), the measurement unit (household vs. individual), and the type of equivalence scale used for adjusting household-level data, amongst other dimensions. As pointed out by Atkinson and Brandolini (2009), it is often not sufficient to account for these differences using dummy variables for each category underlying the data that are being used in a regression, as has been frequently done in the literature. Doing so implicitly assumes that the differences between the types of unit remain constant over time, which has been

Acknowledgements: This essay is joint work with Stephan Klasen. We are grateful to Sebastian Vollmer, Axel Dreher, and the participants of the 2016 Development Economics and Policy Conference, the 2015 AEL Doctoral Research Seminar, the RTG Globalization and Development workshops, the 11th ISI conference on Economic Growth and Development, and the 2015 GlaD conference for valuable comments and suggestions.

shown to not be generally true.[1] More recently, Solt (2016) has, based on the WIID, used imputation techniques to also attempt to address data gaps and consistency issues in his Standardized World Income Inequality Database (SWIID). Although this approach has also been criticized (Jenkins 2015), we will rely on these data in our analysis, but also show the robustness of our results to using the WIID data.

A second focus of the debate was the empirical specification of the inequality-growth relationship. In particular, Forbes (2000) moved from the cross-sectional setting used by Deininger and Squire (1998) to a panel setting for two reasons. First, she wanted to address unobserved heterogeneity through fixed effects (and endogeneity through the use of GMM-type methods). Second, a fixed effects specification which exploits the within-variation is also the more policy-relevant question, as policy-makers are interested in whether changes in inequality in a country will promote or hurt subsequent growth. This approach came at the cost of using rather short panel periods of only five years. Essentially, this time span implies examining the short-term impact of changes in inequality on growth. While interesting, it is not so closely related to the theoretical literature which generally focused on longer-term impacts of inequality on growth (e.g., Galor and Zeira 1993; Alesina and Rodrik 1994. Forbes found that rising inequality is associated with higher subsequent growth, although the result is not significant when 10 year periods are used.

The paper by Forbes attracted a lot of debate and commentary. Apart from the abovementioned data issues (her analysis was based on the DS dataset), there was the concern that the use of fixed effects takes out most of the variation in the dataset and that the little within variation might be heavily affected by measurement error (Knowles 2005). Second, there was concern about the functional form. In particular, Banerjee and Duflo (2003) argued that the data are more consistent with the claim that any change in inequality (whether positive or negative) is associated with lower subsequent growth, which is, of course, a rather different interpretation. There have been further debates on this issue (some of which we address below), but the question how inequality affects growth in a panel setting remains open.

There are three further reasons to revisit this debate again. First, we now have an additional 15 years that can be used to study whether the relationship holds in a longer panel. Second, there have been further improvements in coverage and consistency of inequality data so that one can examine this relationship with an improved data set on inequality. And third, it is important to consider to what extent the relationships found by Forbes (2000) relate to the unique experiences of transition countries. This relates to a separate literature that has pointed out that transition countries experienced a large negative output shock at the start of the transition period in the early 1990s from which they recovered in the late 1990s and early 2000s. More importantly, this initial output shock was associated with a large increase in inequality. In fact, as shown by Ivaschenko (2003) and Klasen and Gruen (2001), the size of the output shock in transition countries was positively correlated with the size of the increase in inequality up until the mid-1990s. The changes in inequality in transition countries in the 1990s and 2000s were among

[1] See, e.g., the examples provided in Atkinson and Brandolinie (2009) on the increasing difference between pre- and post-tax income.

the largest to be found anywhere in the world so that this unique experience, causing a concurrent increase in inequality as well as a collapse and subsequent recovery in economic growth, could potentially be driving the results.

In this paper, we therefore revisit the inequality-growth relationship using an enhanced panel data set with improved inequality data and special attention to the role of transition countries. We base our analysis on the specification of Forbes (2000), but also consider other specifications (including those of Banerjee and Duflo 2003). We find that, when using her specification and the full sample, higher inequality is still significantly associated with higher subsequent growth. But we also find that this finding is entirely driven by the experience of transition countries and is not present in the remaining country sample. It also appears that while increases in inequality are associated with higher growth in transition countries, very rapid and very large increases are associated with reduced growth. However, once we introduce separate time effects for transition countries, these associations disappear as well. We corroborate the finding that there is no robust relationship between inequality and growth in the overall sample with a new instrumental variable strategy introduced by Nunn and Qian (2014) in the literature on aid and growth. We also do not find evidence that the Banerjee and Duflo (2003) specification is superior and cannot confirm symmetry in the relationship of changes in inequality with economic growth.

These results point to three conclusions. First, there is no systematic empirical relationship between initial inequality and growth across the world, except in transition countries. Second, our finding is consistent with the claim that the relationship we find for transition countries is due to the particular timing of inequality and growth dynamics during and after the transition. In particular, the rise in inequality in the 1990s coincided with a sharp output collapse, leading us to find an association between the large increase in inequality in the early 1990 and a growth recovery in the late 1990s. Given that this relationship disappears once separate time effects are introduced for transition countries, it could mean that this association is not causal but rather driven by the particular dynamics of the transition. Instrumenting for inequality, we corroborate the interpretation that there is no systematic relationship between inequality and growth in the transition countries either, and find that the coefficient turns negative, but remains insignificant throughout all tested specifications. Lastly, if the findings for the transition countries are not entirely driven by the temporal dynamics of the transition process, our results suggest that the very low inequality in transition countries at the start of the transition process might have been a barrier to higher growth, but rapid increases apparently were detrimental also there. However, as our instrument is, by construction, not able to pick up the transition experience, we cannot definitely rule out that part of the growth acceleration in the late 1990s originated from the increase in inequality during the transition.

2.2 Literature Review

There is a large theoretical and empirical literature on the relationship between inequality and growth. Because this paper estimates a reduced-form relationship between inequality and growth and does not explicitly test any particular channel through which inequality might affect economic growth, we will not discuss the theoretical literature in detail, but rather give a broad overview of different types of arguments and direct the interested reader to excellent summary papers of the respective field. Following Voitchovsky (2009), theoretical papers can be broadly divided into four types of arguments relating to different parts and aspects of the income distribution.

The first group of papers relates to the circumstances of the poor. One of the arguments most frequently brought forward for why inequality can be bad for growth is that of missed investment opportunities for those at the bottom end of the distribution. Credit market imperfections are the basis for the idea that because the poor are subject to credit constraints, this leads to foregone investment opportunities in physical and human capital, and hence foregone economic growth (e.g., Birdsall 2008, Ghatak and Jiang 2002, Deininger and Squire 1998, Galor and Zeira 1993). Other arguments relating to the bottom end of the income distribution pertain to vicious cycles of poverty and crime (e.g., Chiu and Madden 1998, Josten 2003, and poverty and fertility (e.g., Kremer and Chen 1999).

A second group of arguments, focusing on the size and circumstances of the middle class, argues that domestic demand is a crucial factor determining economic growth, and is typically associated with a (relatively) equal income distribution with few poor (e.g., Foellmi and Zweimüller 2006, Murphy et al. 1989). For a more detailed survey of the demand-side type of arguments, see Erhart (2009). A second well-know channel of how inequality and growth are linked through the circumstances of the middle class is the median voter theorem and related political economy arguments, postulating a negative relationship between inequality and growth Arguably, the higher the inequality in a society, the larger is the gap between mean income and the income of the median voter and the higher his preference for redistribution through taxation, which can reduce incentives and thereby dampen economic growth (e.g., Bertola 1993, Alesina and Rodrik 1994, Perotti 1993). An overview of earlier literature on inequality and public spending can be found in Osberg et al. 2004.

Focusing on the upper part of the income distribution, there are a number of arguments pertaining to the concentration of wealth. One of the most frequently used arguments in favor of having a more unequal distribution of wealth is that the rich can provide the savings necessary for making large investments. This goes back to a model by Kaldor (1955). On the other hand, an unequal distribution of income with high "top" inequality can also be detrimental to growth when it is easier for the elite to capture institutions and extract resources from the economy or to move their capital abroad (see, e.g.,Glaeser et al. 2003.

Finally, the overall distance between individuals in a society also matters for inequality. How far individuals or groups in a society are from each other in economic terms can have important repercussions on growth via the formation of social capital and trust. If very large, the distance between individuals can also have explicit negative consequences for growth via social unrest and the socio-political polarization of society (see, e.g., Keefer and Knack 2002, Easterly 2001).

In sum, while there are arguments in both directions, most of the more recent work favors inequality hurting long-term growth. However, the type and extent of inequality matters. Ferreira et al. (2014), for example, distinguishes between "good" inequality, which rewards effort and leads to better performance (analogous to the "incentive" argument), and "bad" inequality, which wastes human potential (analogous to inequality of opportunity). Besides these conceptual differences, the time frame considered also matters for theoretical predictions of how inequality should affect growth.

In terms of empirical evidence from reduced-form estimations of the effect of inequality on economic growth, we will focus on only the most important contributions given the vast number of empirical studies on the topic. The following overview is based on Neves and Silva (2014). Overall, the evidence on the empirical impact of inequality on growth is mixed and remains controversial. However, a pattern emerges with regards to the results obtained using different empirical specifications. Generally, cross-sectional studies (Alesina and Rodrik 1994, Persson and Tabellini 1994, Clarke 1995, Perotti 1996, and Deininger and Squire 1998 tend to find a negative relationship between inequality and growth, whereas panel analyses yield mostly positive or insignificant results. But Knowles (2005) argues that most evidence on the growth and inequality relationship in cross-sectional studies is derived from inequality data which are not fully comparable. Once the heterogeneity in the underlying income concepts is accounted for, he concludes that there is no remaining relationship between income inequality and growth, but that inequality in expenditure is still negatively correlated with growth. The only cross-sectional study explicitly addressing the endogeneity problem is Easterly (2007), who instruments inequality with a country's "wheat-sugar ratio", which is a function of the fraction of land suitable for growing wheat over the fraction of land suitable for growing sugar cane. The idea is based on the hypothesis by Engerman and Sokoloff (1997) that agricultural endowments predict a country's institutional environment. More specifically, growing sugar cane is more prone to large-scale farming involving slave labor, which leads to higher inequality and extractive institutions, whereas wheat production involves family farming and is associated with the emergence of a middle-class and less inequality. Instead of growth rates,(Easterly, 2007) then shows that higher inequality is associated with lower income levels, as well as worse institutions, and lower education. Most of the cross-sectional results should be viewed with caution because they may contain substantial omitted variable bias, given that any unmeasured factors which are associated with both inequality and growth can be wrongly attributed as an effect of inequality on growth.

Although panel data are not able to perfectly resolve this issue, the possibility of introducing fixed effects allows the removal of at least the time-invariant portion of the omitted variable bias, which is also the main explanation for the divergence in findings between cross-sectional and panel studies. Moreover, it is also more useful from a policy perspective to know what happens to growth if inequality changes within a country, which can be estimated only if the data also contain a time-series dimension. However, apart from the abovementioned data problems which continue to persist in many of the panel data studies using the DS1996 or the WIID data, as well as any remaining concerns about omitted variable bias and endogeneity, panel studies do suffer from another shortcoming: since many of the theoretical effects are likely to have an impact over long periods of time, short-run panels that consider 5 or 10 year periods might be too short to pick up these effects. Nevertheless, we limit the discussion to panel data studies in the following, also because they are more relevant for the empirical set-up of this paper.

The most important study in the context of this paper is Forbes (2000), which we also use as the basis for our own empirical set-up. She finds a small, but positive and significant impact of inequality on subsequent economic growth using 5-year averaged growth rates and the DS1996 dataset. Her sample consists of 45 low- and high-income countries during 1975-95. The application of a difference GMM estimator to deal with the upward bias arising from her dynamic panel structure has, however, been shown to be problematic. Roodman 2009 demonstrates that Forbes' results become insignificant once the econometric issue of overidentification is being addressed, which is something we can confirm in our data as well.

Another widely cited study, Barro 2000 finds, for a samples of 40 to 70 countries and 10-year time periods, that higher inequality leads to lower growth in poor countries and higher growth in rich countries, but there is little overall relationship between income inequality and growth. He refrains from using fixed effects in his preferred specification and points to the exacerbation of measurement error with this approach, but his results from a three-stage least-squares estimation do hold qualitatively in a fixed effects specification, although the latter is only able to capture the contemporaneous relationship between inequality and growth.

Banerjee and Duflo (2003) criticize the functional form assumptions made in previous studies and argue that the growth rate is an inverted U-shaped function of net changes in inequality. They further show how this non-linearity can explain the different findings in previous studies. However, their paper has little to say on the fundamental question of whether inequality is bad for growth. Nevertheless, we test their main empirical specifications on our data as well and find no evidence to support superiority of their empirical (non-linear) set-up over ours.

Deininger and Olinto (2000) focus on asset instead of income inequality in their panel of 60 countries, and find a negative and significant relationship with subsequent growth rates. In addition, they confirm the positive relationship with income inequality as found in previous studies, which continues to hold even when asset inequality is retained in the model.

Ezcurra (2007) looks at annual regional growth across the European Union over the 1993-2002 period and concludes that higher inequality is associated with lower growth, thereby contradicting Barro (2000) wo found that inequality is positively related to growth in rich countries - although the differing results of the two studies could also be due to the different time frames they consider. In sum, results from reduced-form panel studies are heterogeneous and despite the continuous improvement of the inequality data since DS1996, data issues as well as concerns about functional form and appropriate estimation techniques keep being raised in the literature.

2.3 Data and Empirical Strategy

Our estimations are based on a sample of 122 countries over the 1961-2012 period, with a total of 712 observations for the level, and 577 observations for the difference specifications (115 countries). This is much larger data set both in its cross-country as well as its time dimension than those that have been used in the literature. The IV estimates rely on a smaller sample of 92 countries, which does not, however, affect the point estimates of inequality.[2] Unless indicated otherwise, estimations are using 5-year averages of growth as the dependent variable and the beginning-of-period Gini, lagged by one period, as the variable of interest. That is, the first time period is 1961-1965 and the last one is 2011-2012,[3] yielding a total of 12 time periods. Except for the GDP data,[4] which is taken from the Penn World Tables (PWT), Version 8.0 (Feenstra et al. 2015),[5] all control variables are as in Forbes (2000): the price level of investment (also from the PWT) is included, which she uses as a proxy for market distortions, and average years of secondary schooling for the population aged over 25 (taken from the Barro and Lee database, Version 2.0) is added separately for males and females.

To add a causal interpretation to our findings, we use an instrumental variable (IV) estimation employing a recently introduced technique (Nunn and Qian, 2014) which allows us to use Easterly's (2007) wheat-sugar ratio as an instrument for inequality in a panel set-up. We are thereby able to address endogeneity concerns more convincingly with the simultaneous use of both fixed effects and IV-estimation. The idea of the instrument is that the interaction between a cross-sectional variable, varying only between countries, and a time-varying variable, which is the same across all countries, is valid if the level of the respective variable is controlled for. This is generally taken care of for the cross-sectional variable by using a fixed effects estimator, and for the time-varying variable by including year dummies. As mentioned, we use Easterly's (2007) wheat-sugar ratio as the cross-sectional variable and interact it with the oil price, which introduces variation in inequality over time. A higher oil price arguably leads to higher inequality numbers because higher oil prices have a disproportionately larger adverse effect on the poor, who spend a larger share of their budget on staple food items and transport, the prices of both of which increase with the oil price (empirical evidence on the oil price-poverty-inequality link stems mostly from country case studies; see, e.g., Naranpanawa and Bandara (2012) for Sri Lanka, or Essama-Nssah et al. (2007) for South Africa). The estimator then compares the difference in growth in years following a high oil price to years following a low oil price in countries that have a high wheat-to-sugar ratio (low inequality) to countries that have a high wheat-to-sugar ratio (high inequality). The identifying assumption is that the

[2]We are able to reproduce table 2.2 using only the subsample used also in the IV. Results are shown in appendix table 2.A.1.

[3]2011-12 is the only period with less than five years. More recent data was not available at the time of writing.

[4]Forbes used Gross National Income data from the WDI.

[5]In choosing the accounting concept underlying the GDP data for growth rates and levels, we follow the recommendations of the PWT and use the (real) output-based growth rates derived from the national accounts as the dependent variable and the expenditure-based current-price level of GDP as the initial level to capture convergence effects.

effect of the oil price on growth will not (systematically) differ between countries with a high and a low wheat-to-sugar ratio through channels other than inequality. More intuitively, the resulting estimator is similar to a difference-in-differences approach, but with a continuous treatment variable (inequality). The sugar-wheat ratio is correlated with inequality levels (corresponding to the pre-treatment differences in an important observable variable), whereas the oil price is correlated with inequality differences (corresponding to a common time effect in the variable of interest). Through the country fixed effects and the year fixed effects, we also take out the difference of the "baseline" growth rate between countries with high inequality levels and low inequality levels and the time trend they have in common (changes in growth). Remaining changes in growth, taking the baseline differences in inequality levels as well as in growth, and common trends in inequality as well as growth into account, are then attributable to changes in inequality if the exclusion restriction is valid - that is, if changes in the oil price do not systematically affect growth differently in countries with high and countries with low sugar-wheat ratios in a way that is correlated with a country's sugar-wheat ratio after controlling for a number of other potential influencing factors.

Our main measure of inequality, the Gini coefficient of net income, is taken from the Standardized World Income Inequality Database (SWIID) (Solt 2016). One of the main advantages of the SWIID is that the data are strongly balanced, i.e., all missings in the final dataset stem from other control variables. The SWIID is based on the World Income Inequality Database (WIID) (UNU-WIDER 2015) and standardizes the rather heterogeneous and unbalanced database by drawing on several other data sources and multiply imputing values to make the resulting data comparable across countries and over time. The final dataset contains 100 imputations for each data point, allowing the researcher to explicitly account for the uncertainty associated with imputing values by using multiple imputation (mi) estimation. Unless indicated otherwise, estimations employ the "mi: estimate" command as provided by Stata, which yields a single coefficient estimate and its corresponding corrected standard error applying Rubin's rule (Rubin 2004). As opposed to the regression results which exploit all of the 100 imputations, the descriptive statistics and graphs are based on the mean value of the Gini across the 100 imputations. In addition to the overall sample, descriptives are reported separately for transition- and non-transition countries. Our classification of transition countries is based on Gruen and Klasen (2012) and includes 22 post-Communist countries, of which the following 15 are part of our sample: Albania, Armenia, Bulgaria, Czech Republic, Estonia, Hungary, Kyrgyz Republic, Latvia, Lithuania, Poland, Romania, Russia, Slovak Republic, Slovenia, and Ukraine. Table 2.1 contains descriptive statistics for all variables used in the model.

As one can see, most variables do not display major differences between transition- and non-transition countries, notable exceptions being schooling of both males and females, and, very importantly, inequality. The average Gini coefficient in transition countries is a full 8.5 Gini points lower than in non-transition countries, substantiating our belief that the inequality-growth relationship in transition countries is inherently different from that in the rest of the world - or at least the part covered by our sample.

[6] The Price Level of investment (PI) is defined as the PPP over GDP divided by the exchange rate

Table 2.1: Descriptive Statistics

Total sample (712 obs.)	mean	sd	min	max
Gini	38.09	10.52	15.8	75.71
GDP per capita growth	0.023	0.032	-0.199	0.112
Price level of investment[6]	0.65	0.45	0.07	5.93
Initial GDP per capita (in 2005 PPP USD)	11533.3	11414.4	272.8	76523.6
Schooling (female)	2.22	1.56	0.02	6.89
Schooling (male)	2.6	1.52	0.15	7.25
Only transition countries (71 obs.)				
Gini	30.51	6.01	18.87	44.7
GDP per capita growth	0.02	0.054	-0.154	0.112
Price level of investment	0.58	0.21	0.21	1.01
Initial GDP per capita (in 2005 PPP USD)	10936.2	5899.3	1974.7	24519.5
Schooling (female)	3.68	1.15	0.99	6.47
Schooling (male)	3.89	1.03	1.46	6.62
Sample without transition countries (641 obs.)				
Gini	38.98	10.59	15.8	75.71
GDP per capita growth	0.023	0.028	-0.199	0.109
Price level of investment	0.66	0.47	0.07	5.93
Initial GDP per capita (in 2005 PPP USD)	11595.3	11842	272.8	76523.6
Schooling (female)	2.078	1.52	0.02	6.89
Schooling (male)	2.46	1.5	0.15	7.25

All estimations employ country fixed effects to control for unobserved heterogeneity and remove a potential source of (time-invariant) omitted variable bias. While some concerns have been raised in the literature that this approach exacerbates measurement error and removes a large part of the variation in inequality (e.g., Knowles 2005), the use of the more consistent SWIID data, which combine information from different datasets and thereby minimize measurement error, as well as an increase of the within-country variation in inequality in the past 15 years,[7] lead us to believe that these drawbacks no longer justify not using a within estimator. Because of the use of growth rates as the dependent variable and the initial GDP per capita level variable as a control, the fixed effects specifications suffer from Nickell bias, entailing an upward bias on our variable of interest (Nickell 1981). All significant estimates are therefore furthermore subjected to a difference Generalized Method of Moments (GMM) estimator (Arellano and Bond 1991). The estimator eliminates the bias by using deeper lags of the independent variables as instruments, which are by construction uncorrelated with the error term. Orthogonalizing the instruments mitigates the unbalancedness of the dataset. Using the full instrument set would lead to the problem of too many instruments, which in this case exceeds the number of cross-sections (122) and renders the Hansen test of overidentification invalid. In all our

multiplied by 100.

[7]The within-country variation of net income inequality has increased from 14% of the overall variation in the 1960-1996 sample to 18% of the overall variation in the sample going until 2012. While this may still seem rather small, within-country variation of market inequality has increased from 24 to 32%, implying that some of the observed lack of within-variation is the result of successful redistribution.

reported GMM estimates, the instrument set has therefore been restricted in different ways.[8] Because the multiple imputation command does not produce test statistics for the relevant GMM misspecification tests (AR1, AR2, and overidentification tests), they have been conducted individually for each of the 100 imputations. We then report the share of incorrectly specified regressions, along with the mean value of each test statistic. The multiply imputed regressions are considered well specified if less than 5% of the individual regressions are misspecified. In line with Forbes (2000), we use the difference GMM estimator. A system GMM (Blundell and Bond 1998) is sometimes suggested in the literature because the use of the level equation implies that the estimator is less prone to measurement error. However, although the system GMM estimator does yield similar estimates, the results are less clear, and, more importantly, the misspecification tests indicate problems in all but a few cases. System GMM is therefore retained as a robustness check, but the preferred estimator is a (two-step) difference GMM. Standard errors are robust in all estimations as per Windmeijer's (2005) correction procedure.

2.4 Results and Discussion

Table 2.2 displays the first set of results. The first column corresponds to Forbes' (2000) basic specification. Like Forbes, we find a positive coefficient on the inequality variable, although the coefficient is substantially smaller than hers, and, as found by Roodman (2009), this effect does not hold with a non-biased GMM estimator. Appendix table 2.A.2 displays the results for different instrument restrictions, none of which are well specified. Moreover, although the coefficient is now closer to Forbes' estimate of 0.0013, it loses significance in most specifications. Once we include a transition country dummy in column 2 and interact it with the inequality measure, the results become much clearer. The coefficient on the interaction is now substantially larger and highly significant. Moreover, the effect persists in the GMM specification, as shown in column 3. This time, we are also able to find a well-specified regression, which further underpins our belief that the inequality-growth relationship during the transition is inherently different from that in the rest of the sample and that it is incorrect to estimate a common slope parameter for the two processes. Notably, as the transition countries pick up the positive effect of inequality on growth, the non-interacted inequality variable shrinks substantially and turns insignificant. That is, we do not find any effect of inequality on growth in the remaining (non-transition) countries and our findings lead us to conclude that the small positive impact found in the full sample is not robust and is furthermore driven by a small group of transition countries. According to the fixed effects estimate, which is the lowest of our point estimates for the impact of inequality on growth in transition countries, a ten point increase in a country's Gini coefficient - which is roughly equal to the total increase in inequality in transition countries between 1985 and today - would lead to a 4 percent increase in average annual growth over the next five years. However, this result is to be

[8] Instruments have been restricted to a maximum of 2 lags, and collapsed in some cases. The restrictions imposed on the individual GMM regressions are reported in the respective table notes. Our results do not depend on the type of instrument restriction used and we report the ones which perform best in terms of the share of misspecified regressions (which is explained just below).

taken with caution. The processes occurring in the 1990s in transition countries after the breakdown of the Soviet Union - political and economic liberalization, the introduction of market economies and opening up of markets to (non-Soviet) external trade - were exogenous events with effects on both inequality and growth.

Figure 2.1 illustrates the average correlation across all transition countries between inequality and growth as it occurs in the estimation, that is, with the Gini coefficient lagged by one period. A striking image emerges with a sharp increase in both growth and lagged inequality between 1995 and 2000, raising concern that this period might be driving the effect in transition countries. Moreover, it appears as if it is precisely the 5-year lag structure used in our estimations which causes this correlation. Nevertheless, one should be cautious in interpreting the graph since it merely displays the averages across all transition countries, and developments within single countries do not necessarily show the same correlation as depicted here. Indeed, when consulting the individual correlations in each country (as shown in appendix figure 2.A.1), the picture is less clear. An outlier analysis [9] does not yield any clear results pertaining to the issue, either - no single country-year observation is driving the positive impact of inequality on growth in the transition countries.

In order to capture the events occurring in the 1990s which might be driving the observed correlation between inequality and growth at least partially, we introduce separate time effects for the group of transition countries. Indeed, once the separate period dummies are introduced, the positive impact of inequality on growth disappears also for the transition countries, and remains very small and insignificant for the remaining sample (column 4 of table 2.3). Finally, we re-estimate the model of column 1 (corresponding to Forbes' basic specification) in column 5, and introduce separate year dummies for transition countries without including an interaction between the inequality measure and the transition country dummy. The mere introduction of a separate time effect for the transition countries slashes the positive coefficient of inequality by more than half and wipes out the previously found significant positive effect of inequality on growth. In sum, we cannot confirm that higher inequality enhances economic growth in our sample of countries outside the transition period, at least not in terms of higher levels - as opposed to increases or decreases - of inequality.

Building on Banerjee and Duflo (2003), who focus on the relationship between changes in inequality and growth, we also test Forbes' specification in differences instead of levels of inequality.[10] Neither in the full sample, nor using a transition country interaction - with and without separate time effects - do we find any significant impact of changes in inequality and growth.[11]

[9] Instruments have been restricted to a maximum of 2 lags, and collapsed in some cases. The restrictions imposed on the individual GMM regressions are reported in the respective table notes. Our results do not depend on the type of instrument restriction used and we report the ones which perform best in terms of the share of misspecified regressions (which is explained just below).

[10] Note that only the inequality measure has been differenced and the specification does not correspond to a model in differences.

[11] Results available upon request.

Table 2.2: Baseline specifications in levels, FE results

Dep. var.: GDP growth	(1) multiple imputation estimation	(2) transition country interaction	(3) transition country interaction, diff. GMM	(4) transition country interaction & transition-year dummies	(5) transition-year dummies
Gini(t-1)	0.000472*	0.000140	-0.000103	0.000149	0.000169
	(0.000269)	(0.000234)	(0.000612)	(0.000228)	(0.000222)
Transition* Gini(t-1)		0.00400***	0.00653**	0.000565	
		(0.00150)	(0.00282)	(0.00129)	
GDP(t-1)	-0.0513***	-0.0469***	-0.0652***	-0.0420***	-0.0423***
	(0.00895)	(0.00897)	(0.0139)	(0.00890)	(0.00881)
PI(t-1)	-0.00834	-0.00766	-0.00322	-0.00902	-0.00906
	(0.00515)	(0.00527)	(0.0104)	(0.00579)	(0.00579)
Schooling_m(t-1)	2.03e-05	0.00260	0.00107	-0.00205	-0.00207
	(0.00852)	(0.00894)	(0.0175)	(0.00775)	(0.00775)
Schooling_f(t-1)	0.00308	-0.000914	0.000752	0.00155	0.00161
	(0.00922)	(0.00979)	(0.0168)	(0.00952)	(0.00949)
Constant	0.415***	0.382***		0.360***	0.362***
	(0.0694)	(0.0714)		(0.0698)	(0.0689)
# of instruments			74		
AR1			0,0013559		
AR2			0,4196453		
Hansen test			0,1833639		
% misspecified			0		
Observations	712	712	590	712	712
# of countries	122	122	116	122	122
Trans-Year FE	NO	NO	NO	YES	YES
Year FE	YES	YES	YES	YES	YES

Notes. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Instruments in the GMM estimator (column 3) have been orthogonalized and restricted to lags 3 and 4. p-values are reported for the GMM misspecification tests (AR1, AR2, Hansen test). The system GMM estimate can be found in columns 1-3 of appendix table 2.A.2

Figure 2.1: Correlation between growth and (lagged) inequality in transition countries

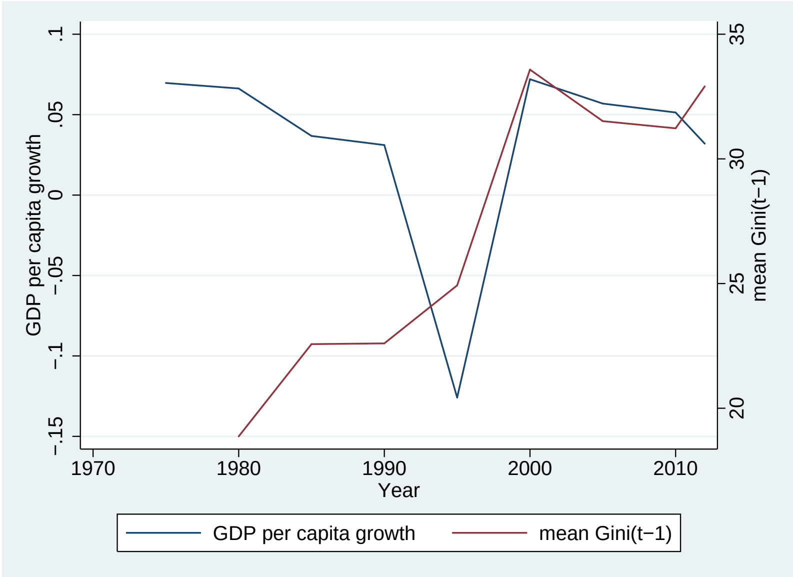

We then introduce both levels and differences simultaneously as shown in table 2.3. In column 1, when both levels and differences are included in the estimation, the positive coefficient of the level of inequality is confirmed, but inequality increases are associated with lower growth. Introducing the transition country interaction in column 2, it becomes clear that these effects are, again, driven by the transition countries. This would, at first glance, suggest that higher inequality levels in transition countries are associated with higher growth, but sharp increases lead to lower growth. However, like in the previous set of regressions, when we proceed to introduce transition-year effects, they eradicate the significance of the coefficients on the transition country-inequality interactions both for levels and differences.[12]

In order to add a causal interpretation to our main findings from table 2.2, we employ an IV strategy using the interaction of the wheat-sugar ratio and the oil price as described above. We use a one-period lag of the oil price to instrument for inequality,[13] and repeat the specifications from columns 1, 2, and 4. The second stage results are presented in the top panel of table 2.4, and the first stage is displayed in the bottom panel, along with

[12] However, the coefficients are not reduced as much as in the equation containing only the levels. When only the time dummies, but not the interaction for transition countries are introduced in column 4, the level effects for the whole sample are significant at the 10% level - however, when this effect is tested in a GMM framework, both the inequality level and change variables lose significance and decrease in size, in line with the direction of Nickell bias. It appears as if, once the effect of changes in inequality is accounted for separately, higher inequality levels are associated with higher subsequent growth in transition economies. Although the coefficient is insignificant, it drives up the size of the effect in the overall sample and, if no separate time dummies are introduced, may lead to interpretations of a substantive inequality-growth relationship in these countries.

[13] There are two reasons for doing so: Firstly, the oil price cannot be expected to have an immediate impact on inequality given that it needs to work its way through the economy and even if it affects production immediately - which is not always plausible - it will not affect prices, e.g. for food, right away. Second, on statistical grounds, the first lag turns out to be the stronger instrument, although the contemporaneous oil price is also valid and delivers very similar estimates (results available upon request). When both the contemporaneous variable and the first-period lag are included as instruments, only the lagged oil price turns out to be significant. We therefore decided to drop the contemporaneous variable from the instrument set.

Table 2.3: Baseline specification, augmented with differences

Dep. var.: GDP growth	(1) FE	(2) FE	(3) FE	(4) FE	(5) diff. GMM
Gini(t-1)	**0.000947****	**0.000518**	**0.000532**	**0.000577***	0.000377
	(0.000384)	(0.000338)	(0.000326)	(0.000316)	(0.000440)
Transition* Gini(t-1)		**0.00426*****	**0.00145**		
		(0.00155)	(0.00118)		
Δ Gini(t-1)	**-0.000653****	**-0.000370**	**-0.000365**	**-0.000442***	-0.000204
	(0.000288)	(0.000262)	(0.000257)	(0.000256)	(0.000430)
Transition* Δ Gini(t-1)		**-0.00236****	**-0.00143**		
		(0.000919)	(0.00119)		
Observations	577	577	577	577	577
#of countries	115	115	115	115	115
Control variables	YES	YES	YES	YES	YES
Trans-Year FE	NO	NO	NO	YES	YES
Year FE	YES	YES	YES	YES	YES

Notes. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Instruments in the GMM estimator (column 7) have been restricted to lag 3, resulting in a total of 97 instruments. Misspecification tests (p-values) of the GMM are: AR1: 0.0058, AR2: 0.3473, Hansen test: 0.3473, misspecified: 0%. The system GMM estimate can be found in columns 4-6 of appendix table 2.A.3

weak-instrument tests. Our instruments turn out to be valid and are clearly above the cut-off value in the first two specifications, and still provide reasonable estimates in the third one.[14]

Inequality is not only insignificant, but the coefficient has now turned negative, and remains so across all four specifications. This result is in line with the GMM finding from table 2.3, and indicates that there exists some form of endogeneity introducing a spurious positive correlation between inequality and growth in the OLS estimates. What is more, the instrument appears to work even with the transition country interaction. Although separately instrumenting for transition countries takes out a good deal of variation from the instrument given that they all are Eastern European and Central Asian countries and have positive wheat-sugar ratios, the IV still appears to deliver reasonable estimates. Since the IV is biased towards the OLS estimate in the presence of weak instruments, concern with weak instruments does not threaten the validity of our conclusions from column 3. If anything, the true parameter estimates might be even further from the positive association found in the OLS/FE specifications.

Naturally, the instrument does not pick up the changes in inequality caused by the transition and, even with the introduction of separate transition-period dummies in column 2, remains to be a strong instrument for inequality. Overall, these results substantiate our belief that the positive association between inequality and growth is non-causal, although

[14]This is the F-statistic for the single-instrument case (columns 1 and 2) and the maximum bias in percent according to the Kleibergen-Paap test statistics, which is appropriate when several instruments and robust standard errors are used (column 3). Our instrument turns out to be strong as per the cut-off value of 10 for the first-stage F-statistics, and while the coefficient in column 3 may be biased up to 10% towards the OLS estimate, this does not cause a problem for the results here given that the inequality variables are insignificant and very different from the OLS/FE results.

Table 2.4: Baseline specifications, IV results

Dep. var.: GDP growth	(1) basic	(2) transition-year effects	(3) transition country interaction
Gini(t-1)	-0,00221	-0,000688	-0,000735
	(-0,0014)	(-0,00135)	(-0,0016)
Transition* Gini(t-1)			-0,00725
			-0,00593
GDP(t-1)	-0.0503***	-0.0415***	-0.0649***
	(-0,0118)	(-0,00922)	(-0,0176)
PI(t-1)	-0,00321	-0,0055	-0,0051
	(-0,00594)	(-0,0056)	(-0,0058)
Schooling_m(t-1)	-0,00851	-0,00443	-0,0129
	(-0,00865)	(-0,00654)	(-0,0106)
Schooling_f(t-1)	0,0076	0,000866	0,0165
	(-0,00911)	(-0,00718)	-(0,0123)
Observations	566	566	566
# of countries	92	92	92
Year FE	YES	YES	YES
Transition-Year FE	NO	YES	NO
FIRST STAGE			
GDP(t-1)	2.505*	3.554***	2.477*
	(-1,318)	(-1,268)	(-1,327)
PI(t-1)	0,774	0,824	0,782
	(-0,701)	(-0,683)	(-0,703)
Schooling_m(t-1)	-2.709*	-2,092	-2.664*
	(-1,401)	(-1,386)	(-1,383)
Schooling_f(t-1)	1,499	0,535	1,464
	(-1,546)	(-1,534)	(-1,53)
SWratio*Oil price(t-2)	0.315***	0.312***	0.303***
	(-0,0835)	(-0,091)	(-0,093)
SWratio*Oil price(t-2)*trans			0,0455
			(-0,153)
R^2	0,127	0,192	0,127
Weak instruments: F-stat	14,2	11,79	
Kleibergen-Paap max. bias			10%

Notes. Robust standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1. Note that estimation is not based on multiple imputation because the combination of mi:estimate with xtivreg2 is currently not yet implemented in Stata. We use the average Gini across the 100 imputations instead. The point estimate using the average estimate across the 100 imputed data sets yields a coefficient of -0.00224, which is very close to the one estimated above.

we cannot definitely rule out that the increase from excessively low inequality levels during the transition did contribute to the subsequent high growth rates in the transition countries as well.

One should keep in mind, however, that even after the transition, inequality in transition countries is still rather low compared to non-transition countries: the maximum inequality value found among the transition countries is still only around half a standard deviation above the non-transition country mean. Our tentative reading of the results on the transition countries thus far could be that the higher inequality levels in these countries after the transition might therefore rather represent "normal" inequality levels, and the previous, excessively low inequality numbers could reflect the fact that inequality was kept at "artificially" low levels due to the income compression during the socialist system. The breakdown of the Soviet Union led to a well-documented (see, e.g., Aristei and Perugini 2012) large drop in output with a recovery over the 1990s, and the transition to a market economy was associated with unprecedented increases in inequality. Although we do not specifically test for this, our reading of the literature on the topic is that the two developments in inequality and growth are more likely to have been unrelated events. That is, both are caused by the transition, but one is not causing the other, although the possibility cannot be definitely ruled out on the basis of our estimations.

2.5 Robustness Tests

As a check on functional form, we test Banerjee and Duflo's proposition that changes in inequality may just be measurement error, and because measurement error is larger in times of economic distress, this would cause a negative relationship between changes in inequality and growth. Despite the fact that their argument would entail a contemporaneous relationship between inequality changes and growth and we are estimating a lagged one, we run a number of different specifications to see whether we find a symmetric effect of changes in inequality on growth. If positive and negative changes in inequality are symmetrically offsetting each other, this would also explain why we do not find any effect in the difference equations. In order to generally account for functional form issues brought up by Banerjee and Duflo, we also test the level equation for such effects. In a first step, we are simply including a quadratic term in both the level and the difference specifications. Table 2.5 displays the results.

The only significant result is that of the difference specification with the transition country interaction (column 2). An F-test of joint significance indicates that the effect is significant at the 1% level. At a value of 141.7, the maximum is located far from even the highest of the transition country Gini coefficients of 48.5, and even further from the mean of 29 Gini points. The result would hence indicate that the sample values are located on the upward-sloped part of the curve, meaning that positive changes in inequality enhance growth, but at a decreasing rate. However, when subjected to a difference GMM, none of the quadratic terms were jointly significant (as shown in table 2.A.3). We therefore reject the proposition of a quadratic effect of inequality on growth for transition countries as well as non-transition countries, and in both levels and differences.

Table 2.5: Quadratic FE specifications in differences

Dep. var.: GDP growth	(1)	(2)	(3)	(4)
Gini(t-1)	0,00205	-0,000164		
	(-0,00156)	(-0,00123)		
Gini(t-1)^2	-1,76E-05	3,29E-06		
	(-1,56E-05	(-1,24E-05)		
Transition*Gini(t-1)		0,00717		
		(-0,012)		
Transition*Gini(t-1)^2		-5,06E-05		
		(-0,000183)		
ΔGini(t-1)			-0,000209	-0,00013
			(-0,000216)	(-0,000212)
ΔGini(t-1)^2			-1,09E-05	-1,09E-05
			(-1,33E-05)	(-1,17E-05)
Transition*ΔGini(t-1)				-0,00141
				(-0,00119)
Transition*ΔGini(t-1)^2				7,58E-05
				(-8,42E-05)
Observations	712	712	577	577
#of countries	122	122	115	115
Control variables	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Turning point for trans.	58.2 (Max.)	**141.7 (Max.)**	No quadr. effect	9.3 (Min.)
F-test of quadratic terms	0,2363	**0.0393*** **		0,4941

Notes. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. The p-value is reported for the F-test. The turning point of the quadratic effect is only calculated for the transition countries, i.e., from the coefficients Transition*l.Gini and Transition*l.Gini^2 (and their respective values in differences. The F-test of quadratic terms test contains all constituent terms of the interactions, i.e., L.Gini, L.Gini^2, Transition*l.Gini and Transition*l.Gini^2

Piecewise linear regressions

As another test of the functional form concerns, we run a set of piecewise linear regressions. They are based on inequality changes, and employ different margins of change ranging from 3 to 20% change in inequality, as indicated in the top row. Differential slopes are estimated for negative, zero (within the aforementioned margin), and positive changes. This is similar to Banerjee and Duflo's (2003) piecewise linear approach, but instead of using the model in differences, we are basing the inequality change brackets on the changes in levels of inequality since no evidence for any kind of relationship between inequality and growth was found in the differenced specification in the first step of our analysis (appendix table 2.A.5).[15] The FE estimates (table 2.6, columns 1-4) show that growing inequality is related to lower subsequent growth.

[15]The relationship is estimated with and without including the level variable of the Gini into the model, but results are almost identical between the two specifications (this is true for all versions of the piecewise linear specification, including the subsequent versions using subsamples and interactions) and we therefore proceed with the model without the level variable. Appendix table 2.A.6 displays the results with the level variable.

Table 2.6: Piecewise linear regressions of inequality changes, FE and GMM results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. var.:	FE	FE	FE	FE	GMM	GMM	GMM	GMM
GDP growth	3	5	10	20	3	5	10	20
Neg. change	-0,000217	-0,000114	5,41E-05	0,000513	0,00111	0,00113	0,00131	0,00194
	(0.000283)	(0.000323)	(0.000443)	(0.000858)	(0.000788)	(0.000813)	(0.000922)	(0.00131)
No change	0.00100*	0,000367	2,92E-05	-0,000126	-0,000221	-8,09E-05	-0,000106	3,92E-05
	(0.000509)	(0.000337)	(0.000200)	(0.000142)	(0.00128)	(0.000712)	(0.000416)	(0.000331)
Pos. change	-0.000995***	-0.00103***	-0.00116***	-0.00141***	-0.00164***	-0.00172***	-0.00197***	-0.00264***
	(0.000213)	(0.000225)	(0.000254)	(0.000319)	(0.000535)	(0.000585)	(0.000600)	(0.000949)
# of Instr.					93	93	93	93
AR1					0,0012974	0,0015662	0,0015634	0,0031463
AR2					0,9706342	0,9785592	0,9778599	0,9689
Hansen test					0,436301	0,4110748	0,3200502	0,266292
% misspcfd.					0	0	0	0
Observations	614	614	614	614	497	497	497	497
# of countries	115	115	115	115	110	110	110	110
Control vars.	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES	YES	YES

Notes. Robust standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.Instruments in the difference GMM have been restricted to lags 3 and 4. The results with the restricted instead of the collapsed instrument set are reported here due to problems with the misspecification test for the 10%-change specification (column 7). Using collapsed instruments, the results are very similar for the positive changes, but negative changes are also significant (results available upon request). Numbers in the third row represent the knots for defining the "no change"-bracket, i.e., changes between $+/-3$ (5, 10, 20) percent are coded as "no change", and changes above (below) the knot as increases (decreases).

This relationship is confirmed in both the difference- and the system GMM estimations. The association is stronger, but less robust for larger changes in inequality. No robust relationship is found for negative inequality changes, but the coefficients are mostly positive, especially for the larger changes, and are significant in some of the GMM specifications. When the same estimation is repeated with a subsample excluding transition countries, the coefficients on the positive change variable retain their negative sign, but become insignificant. The results can be found in appendix table 2.A.7.[16]

Finally, the specification using the full sample is repeated, but with interactions between the inequality change variables and a transition country dummy (table2.7). Although some of the positive coefficients are insignificant in the GMM estimations, the results clearly show that the negative and significant effect of positive inequality changes on growth stems from the transition countries only. In line with the results using only the subsample of non-transition countries, the coefficient on the positive change variable remains negative, but it is very small and far from significant. Again, once the transition country dynamics are accounted for separately (columns 5-8), no significant impact of inequality is found for the remaining sample. The results of the piecewise linear regressions indicate that when a separate slope is estimated only for those countries showing positive inequality changes, higher inequality increases are actually associated with lower growth rates. This finding directly contradicts the proposition put forward by Banerjee and Duflo (2003) that changes in inequality could just be measurement error, which would imply a symmetric effect of both positive and negative inequality changes being associated with lower growth. However, the positive effect is, again, driven by the group of transition countries and is not robust to transition-year effects.

[16]Because the FE results are insignificant, they are not further subjected to a GMM.

Table 2.7: Piecewise linear regressions of inequality changes with transition country interaction, FE results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. var.: GDP growth	3	5	10	20	3	5	10	20
Decrease	-0,000365	-0,000274	-0,000166	0,00022	-0,000351	-0,000249	-0,000136	0,000218
	(0.000276)	(0.000318)	(0.000442)	(0.000820)	(0.000273)	(0.000314)	(0.000434)	(0.000814)
Decrease*trans	0,00112	0,00122	0,00121	0,00167	-7,35E-05	-0,000431	-0,000775	-0.00629**
	(0.00136)	(0.00150)	(0.00201)	(0.00307)	(0.000999)	(0.00111)	(0.00180)	(0.00296)
No change	0,000317	-1,34E-05	-0,000128	-0,000143	0,000162	-0,000127	-0,000196	-0,000188
	(0.000458)	(0.000309)	(0.000189)	(0.000136)	(0.000457)	(0.000309)	(0.000188)	(0.000135)
No change*trans	-4,33E-05	-0,000101	-6,30E-05	-0,000596	0,000737	0,00094	0,000725	0,000662
	(0.00161)	(0.00109)	(0.000752)	(0.000672)	(0.00137)	(0.000879)	(0.000615)	(0.000426)
Increase	**-0,000151**	**-0,000133**)	**-0,000132**)	**-0,000188**)	**-0,000149**)	**-0,000126**)	**-0,000118**)	**-0,000165**)
	(0.000223)	(0.000240)	(0.000294)	(0.000416)	(0.000220)	(0.000238)	(0.000294)	(0.000414)
Increase*trans	**-0.00141*** **	**-0.00144*** **	**-0.00150*** **	**-0.00134** **	**-0,000172**	**-0,000219**	**-0,000292**	**-0,000272**
	(0.000379)	(0.000391)	(0.000448)	(0.000582)	(0.000380)	(0.000388)	(0.000429)	(0.000541)
Observations	614	614	614	614	497	497	497	497
# of countries	115	115	115	115	110	110	110	110
Control variables	YES	YES	YES	YES	YES	YES	YES	YES
Trans-Year FE	**NO**	**NO**	**NO**	**NO**	**YES**	**YES**	**YES**	**YES**
Year FE	YES	YES	YES	YES	YES	YES	YES	YES

Notes. Robust standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1. Instruments in the difference GMM have been collapsed. Numbers in the second row represent the knots for defining the "no change"-bracket, i.e., changes between +/−3 (5, 10, 20) percent are coded as "no change", and changes above (below) as increases (decreases).

Alternative time spans and lag structures

We also test our main specification for robustness to the choice of the lag structure as well as the time span chosen. Forbes (2000) also included 10 year averages in her analysis and found in what she called an "informal test" that the positive relationship between inequality and growth diminished over time, but noted that because of the limited degrees of freedom, these results were to be interpreted with caution. Now that we have four new time periods available for estimation, we are repeating the exercise to see whether there are different dynamics for ten- as opposed to five-year periods, and to test whether these effects are equally sensitive to how transition countries are accounted for in the estimation. As shown in table 2.8, the results using ten-year averages do not only qualitatively resemble the 5-year ones, but also the magnitude of the effects is rather similar. This is in stark contrast to Forbes' results, where the $10-$year coefficient on inequality was only little over one third of the $5-$year one. We can also confirm that the same caveats pertaining to the 5-year results are present in the 10-year averaged data as well: the inclusion of transition countries diminishes the positive impact of inequality on growth and renders the coefficient insignificant. Transition countries appear to have a positive relationship between inequality and growth, but once the transition-year effects are included as well (columns 3 and 4), there is no significant association between inequality and growth in neither the transition countries nor the remaining sample.

Table 2.8: $10-$year averages

	(1)	(2)	(3)	(4)
Dep. var.: GDP growth	baseline	transition country interaction	transition country interaction & trans.-year effects	transition-year effects
Gini(t-1)	0.000377*	0,000253	0,000247	0,000268
	(-0,000225)	(-0,000218)	(-0,000216)	(-0,000212)
Gini(t-1)*trans		0.00255***	0,000883	
		(-0,000818)	(-0,000817)	
Observations	296	296	296	296
# of countries	118	118	118	118
Control variables	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Trans-Year FE	NO	NO	YES	YES

Notes. Robust standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

A second concern pertaining to timing is the lag structure. The graphical depiction of the inequality and growth variables in transition countries raises concerns that it is merely the choice of a one period lag which generates the correlation between the two variables. We therefore re-run the basic specification of table 2.2, once with a contemporary time structure and once with a two period lag. The contemporaneous specification, shown in the first panel of table 2.9, does not yield any significant results - if anything, there appears to be a negative contemporaneous correlation between inequality and growth in transition countries, but the effect is not robust to the inclusion of the transition-year effects (column 3). The coefficient on the remaining sample is very small and insignificant throughout. In sum, there seems to be no systematic contemporaneous relationship between inequality and growth, either. More results emerge with the two period lagged Gini coefficient, displayed in the second panel of table 2.9. The estimates are similar to those obtained for the one period lag (including the changes occurring when transition countries and transition-year effects are introduced) but are larger and more significant. Importantly, the coefficient for the overall sample remains positive and significant throughout the fixed effects specifications (columns 4-8).[17] However, when subjected to a GMM,[18] it loses significance as well. We are therefore confident that our results are neither contingent on the choice of a particular lag structure, nor on the use of 5-year averages rather than a longer time span.

[17]We have ruled out that this is simply a sample composition effect. Results using a constant sample from the two period lag specification are available upon request.

[18]Note that the GMM is not based on a multiple imputation estimation due to problems with keeping the sample constant when deeper lags are involved. The corresponding FE estimate (replicating column 7), along with further GMM specifications using other restrictions on the lags can be found in appendix table 2.A.9. Because the non-mi FE estimate is slightly larger and more significant than the one using proper mi estimation, the corresponding GMM estimate is a rather optimistic estimate of the impact of inequality on growth, and the mi estimate can be expected to be slightly lower.

Table 2.9: Alternative lag structures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lags	0	0	0	2	2	2	2	2
Dep. var.: GDP growth	FE	FE	FE, trans. country interaction & trans.-year effects	FE	FE	FE, trans. country interaction & trans.-year effects	FE, trans.-year dummies	GMM, trans.-year dummies
L.Gini	-5,66E-05	0,000104	4,29E-05	0.000793***	0.000518**	0.000515**	0.000560**	0,0000169
	(0.000305)	(0.000301)	(0.000290)	(0.000278)	(0.000256)	(0.000251)	(0.000245)	(0.000785)
L.Gini*trans		-0.00441*	-0,00145		0.00269***	0,0012		
		(0.00226)	(0.00188)		(0.000854)	(0.00119)		
Observations	721	721	721	625	625	625	625	506
# of countries	122	122	122	119	119	119	119	114
Control vars.	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES	YES	YES
Trans-Year FE	NO	NO	YES	NO	NO	YES	YES	YES

Notes. Robust standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1. Instruments in the difference GMM have been collapsed. Numbers in the second row refer to the lag length of the inequality variable.

Alternative inequality data

Although there are some clear advantages to using Solt's 2016 SWIID data, some researchers have expressed concern over the choice of the imputation procedure, and the validity of the resulting data (Jenkins 2015). We therefore repeat our analysis with the WIID data, a previous version of which Forbes' 2000 analysis was also based on. Due to the heterogeneity of the underlying data, most authors use some sort of adjustment to make the Gini coefficients contained in the dataset more comparable (such as adding the average difference of 6.6 Gini points between the expenditure and income based Gini coefficients onto the expenditure one). We use a more sophisticated, regression-based adjustment procedure, based on Gruen and Klasen (2012).[19] Again, as shown in table 2.10, the results are similar to what we have obtained in our basic specifications in table 2.2: the positive and significant coefficient of inequality is driven by the transition countries and vanishes when the transition-year effects are introduced in the estimation (columns 3 and 4), although the coefficient on the interaction just misses significance in column 2.

Table 2.10: WIID (adjusted) Ginis, FE results

Dep. var.: GDP growth	(1)	(2)	(3)	(4)
Gini(t-1)	0.000688**	0,000352	0,000415	0,000322
	(-0,000339)	(-0,000353)	(-0,000345)	(-0,000328)
Gini(t-1)*trans		0,002	-0,000944	
		(-0,00129)	(-0,00106)	
Observations	562	562	562	562
R^2	0,326	0,34	0,483	0,481
# of countries	118	118	118	118
Control variables	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Trans-Year FE	**NO**	**NO**	**YES**	**YES**

Notes. Robust standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1.

Further IV robustness checks

To account for the possibility that the exclusion restriction of the IV might be violated due to the differential production structure of agriculture in economies with high and low wheat-sugar ratios (e.g., wheat producing countries might rely more on machinery and hence could be more dependent on fuel as an input factor and also be hit more adversely by increases in the oil price in terms of economic growth), we include a number of control variables capturing countries' agricultural production structure: the size of land under

[19]The adjustment procedure regresses the full sample of Gini coefficients on the different income definitions and reference units used in the dataset to remove the effect of the differential concepts underlying the data, which are added or subtracted from the reported Gini to achieve at a measure equivalent to that based on gross income per person. Because the resulting dataset contains duplicate observations whenever more than one income concept was available in the original data, we report another version of table 2.9 in the appendix (table 2.A.10), where the duplicates where switched. The results are very similar between the two versions.

cereal production, the use of machinery in agriculture, and total agricultural land (all from the WDI). As shown in table 2.11, none of these variables threaten the validity of our IV. If anything, controlling for the share of land under cereal production leads to a more precise estimate of the inequality coefficient. Note that the change in the sign of the coefficient when agricultural machinery is included is entirely attributable to the smaller sample: running the IV without the variable on the same sample (column 4) yields virtually the same coefficient estimate. We also interact the variables with our instrument to model more explicitly that the instrument might affect countries with different agricultural endowments differently. The results (presented in appendix table 2.A.12) do not indicate that this is the case, with the coefficients on the inequality variable again remaining virtually unchanged. As another check on the robustness of the IV, we have included separate time trends for the OPEC countries to account for the fact that the effect of a higher oil price might affect inequality differently in these countries. While the idea for using the oil price as a correlate of inequality was mainly through the adverse effect of a higher oil price on the poor, it might actually affect inequality through the other end of the distribution in oil producing countries and could thereby have a differential impact on growth in these countries. Including separate time trends for the OPEC countries (results shown in appendix table 2.A.11) does not affect the IV estimates much and merely leads to slightly lower (but still valid) F-statistics in the first stage, which is in line with the reasoning of a different transmission mechanism of the oil price on inequality in OPEC countries. We have also tested the impact of separately including a time trend for each continent. Apart from the "Europe and Central Asia" dummy (again, capturing the transition economies), none of these have a major impact on the estimates.[20]

Table 2.11: Robustness of the IV result to further control variables

Dep. var.: GDP growth	(1)	(2)	(3)	(4)
Gini(t-1)	-0,00213	0,000553	0,000572	-0.00294*
	(-0,00155)	(-0,00271)	(-0,00267)	(-0,00176)
Agricultural land	9,35E-05			
	(-0,000643)			
Agricultural machinery		1,67E-06		
		(-6,48E-06)		
Land under cereal production				-2.83e-09*
				(-1,66E-09)
Observations	563	327	327	563
# of countries	92	75	75	92
Control variables	YES	YES	YES	YES
Year FE	YES	YES	YES	YES
Transition-Year FE	NO	NO	NO	NO

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1.

[20]Results available upon request

2.6 Conclusion

In this paper, we have revisited the inequality-growth relationship using an enhanced panel data set with improved inequality data and special attention to the role of transition countries. We based our analysis on the specification of Forbes (2000), but also addressed the functional form concerns raised by Banerjee and Duflo (2003). Using the SWIID data, which provide an improved and substantially longer panel dataset, we can avoid several of the data concerns brought up by the literature, such as consistency over time and between countries, and a low within-country variation. We also take into account the unique experience of transition countries, which suffered a large negative output shock at the start of the transition period in the early 1990s from which they slowly recovered in the late 1990s and early 2000s. This was coincidental with large increases in inequality, which had been kept at low levels during the Communist rule.

Using robust dynamic panel estimation and multiple imputation estimation, we find no robust, systematic relationship between inequality and subsequent growth, neither for levels nor for changes in inequality. While higher inequality appears to be significantly associated with higher subsequent growth when Forbes' and Banerjee and Duflo's basic specifications are used, we find that this effect is entirely driven by the experience of transition countries and is not present in the remaining country sample. Once we introduce separate time effects for the transition countries, these associations disappear for this group of countries as well. These results hold for different lag structures as well as for the medium- rather than the short term, and the empirical patterns observed emerge not only in the SWIID, but also the WIID data.

Our results point to two conclusions. First, there does not appear to be a trade-off between inequality and growth. Second, because the positive impact of inequality on growth in transition countries is not robust to the inclusion of separate time effects, it appears to be driven by other events. Our findings are hence consistent with the claim that the relationship is due to the particular timing of inequality and growth dynamics in transition countries. In particular, the rise in inequality in the 1990s coincided with a sharp output collapse, leading us to find an association between the large increase in inequality in the early 1990 and a growth recovery in the late 1990s.

Results from an IV estimation confirm our interpretation of the positive association between inequality and growth found in the FE specifications as non-causal, both within as well as outside of the transition countries. Given that our instrument does not pick up the transition experience itself, we cannot, however, infer from our IV estimates that the observed positive association between inequality and growth during and after the transition is entirely spurious. However, while we may not be able to definitely rule out a causal link between the increase in inequality and the subsequent growth spell on the basis of our estimations, research on the dynamics of the transition (see, e.g., Aristei and Perugini 2012, and the references in Sukiassyan 2007) suggests that the breakdown of the Soviet regime and the economic transition to a capitalist system triggered both the increase in inequality as well as the slump and subsequent recovery in growth, rather than one causing the other.

2.A Appendix

Figure 2.A.1: Correlation between growth and (lagged) inequality in transition countries

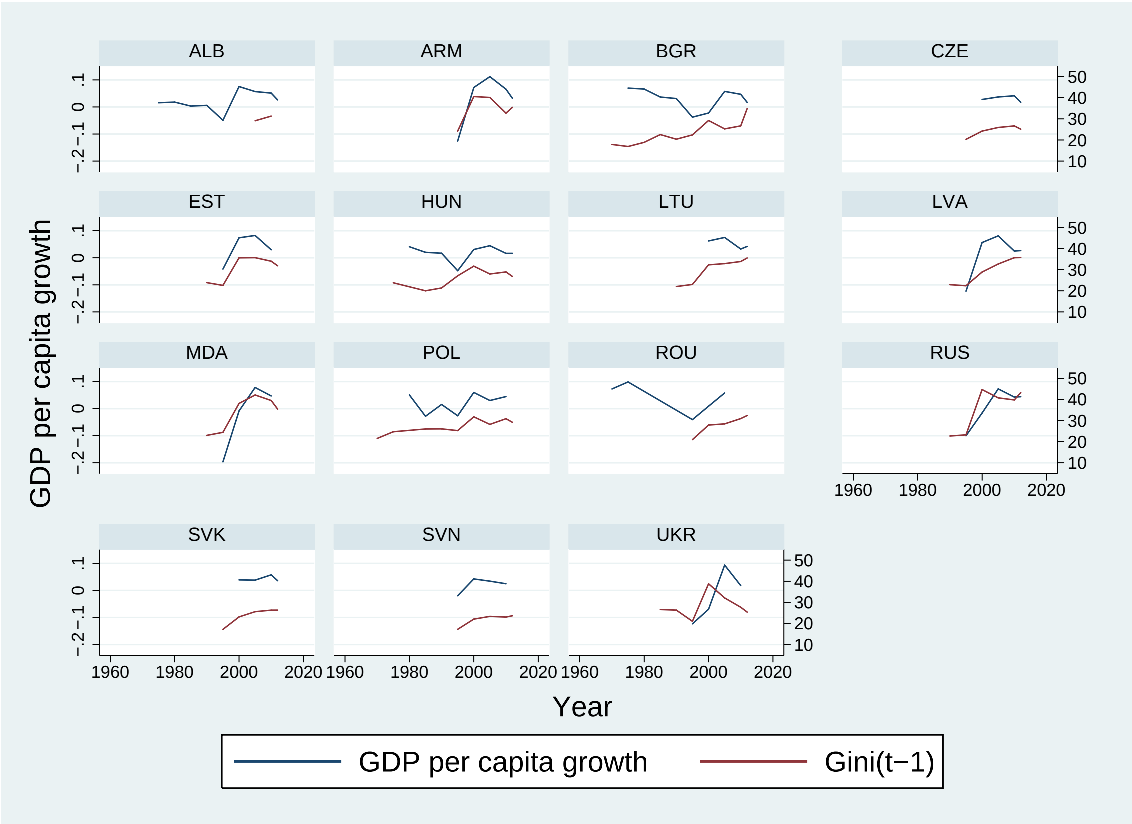

Table 2.A.1: Reproduction of table 2.2 using the IV sample

Dep. var.: GDP growth	(1) levels	(2) levels with transition countries	(3) levels with transition countries & separate transition-year effects	(4) levels with transition-year effects
Gini(t-1)	0.000802**	0.000318	0.000331	0.000379
	(-0.000327)	(-0.000259)	(-0.000252)	(-0.000249)
Transition*Gini(t-1)		0.00463***	0.00112	
		(-0.00175)	(-0.00168)	
GDP(t-1)	-0.0537***	-0.0464***	-0.0404***	-0.0412***
	(-0.0107)	(-0.00945)	(-0.00847)	(-0.00841)
PI(t-1)	-0.00663	-0.00579	-0.00711	-0.00717
	(-0.0049)	(-0.00499)	(-0.00546)	(-0.00547)
Schooling_m(t-1)	0.00342	0.00697	0.0012	0.00111
	(-0.00763)	(-0.00799)	(-0.00621)	(-0.00625)
Schooling_f(t-1)	-0.0027	-0.00857	-0.00479	-0.00456
	(-0.00822)	(-0.00847)	(-0.00762)	(-0.00761)
Constant	0.429***	0.376***	0.346***	0.353***
	(-0.0808)	(-0.0746)	(-0.0658)	(-0.0648)
Observations	587	587	587	587
# of countries	92	92	92	92
Year FE	YES	YES	YES	YES
Transition-Year FE	NO	NO	YES	YES

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.2: GMM results, level specification

Dep. var.: GDP growth	(1) restricted	(2) collapsed	(3) col. & res.	(4) ort. & res.	(5) ort. & col.
Gini(t-1)	**0.00163**	**0.00161***	**0.00178**	**0.00169***	**0.00214**
	(-0.00158)	(-0.000957)	(-0.00504)	(-0.000901)	(-0.00141)
PI(t-1)	-0.036	-0.028	-0.0459	-0.0158	-0.0505**
	(-0.0262)	(-0.0181)	(-0.0798)	(-0.0174)	(-0.0243)
GDP(t-1)	-0.131***	-0.111***	-0.208***	-0.0928***	-0.127***
	(-0.0237)	(-0.021)	(-0.0447)	(-0.0181)	(-0.0192)
Schooling_m(t-1)	-0.0308	-0.0402	0.0133	-0.00455	-0.0282
	(-0.03)	(-0.0288)	(-0.123)	(-0.0203)	(-0.0305)
Schooling_f(t-1)	0.0257	0.0392*	0.0334	0.0202	0.04
	(-0.0306)	(-0.0207)	(-0.0777)	(-0.0198)	(-0.0337)
# of instruments	74	44	19	74	44
AR1	0.0198631	0.13865757	0.8283414	0.00442597	0.0572844
AR2	0.65305763	0.71841735	0.8001032	0.64667875	0.650463
Hansen test	0.05399719	0.00620502	0.3936974	0.1044502	0.0204975
% misspecified	100	100	100	49	100
Observations	566	566	566	590	590
# of countries	115	115	115	116	116
Year FE	YES	YES	YES	YES	YES

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1. The top row indicates the type of instrument restriction (res=lags restricted to 3&4, col=collapsed, ort=orthogonalized).

Table 2.A.3: System GMM results of the basic specifications

Dep. var.: GDP growth	(1)	(2)	(3)	(4)	(5)	(6)
Corresponding GMM specification:	Table 2, column 3	Table 2, column 3	Table 2, column 3	Table 3, column 5	Table 3, column 5	Table 3, column 5
Instrument restriction:	res.	ort. & col.	ort. & res	restricted	ort. & col.	ort. & res.
Gini(t-1)	0.000449	0.0007	0.000429	-0.000387	-0.000374	-0.000303
	(0.000391)	(0.000828)	(0.000331)	(0.000309)	(0.000425)	(0.000436)
Transition*Gini(t-1)				0.00241***	0.00181*	0.00191
				(0.000686)	(0.00110)	(0.00132)
Δ Gini(t-1)	9.66E-05	0.000121	4.03E-05	0.000486	-0.000558	-0.000674
	(0.000274)	(0.000418)	(0.000220)	(0.000469)	(0.000486)	(0.000565)
Transition* Δ Gini(t-1)				-0.00154	-0.000125	0.000449
				(0.00158)	(0.00198)	(0.00260)
GDP(t-1)	0.00423	0.00423	0.00228	0.00184	0.00215	0.0072
	(0.00306)	(0.00306)	(0.00215)	(0.00354)	(0.00522)	(0.00637)
PI(t-1)	-0.0112	-0.0112	-0.0102	-0.0156*	-0.0154	-0.0209
	(0.00736)	(0.00736)	(0.00709)	(0.00902)	(0.0132)	(0.0155)
Schooling_m(t-1)	0.000969	0.000969	0.00322	-0.00922	-0.00801	-0.00904
	(0.00664)	(0.00664)	(0.00570)	(0.00705)	(0.00779)	(0.00594)
Schooling_f(t-1)	-0.00182	-0.00182	-0.00267	0.0059	0.00599	0.00422
	(0.00680)	(0.00680)	(0.00537)	(0.00676)	(0.00776)	(0.00673)
Constant	-0.0122	-0.0122	0.00145	0.0606*	0.0486	0.0135
	(0.0350)	(0.0350)	(0.0292)	(0.0353)	(0.0482)	(0.0549)
Observations	712	712	712	577	577	577
# of countries	122	122	122	115	115	115
Year FE	YES	YES	YES	YES	YES	YES
# of instruments				113	89	59
Hansen Test				0.896	0.592	0.943
AR(1)				0.0113	0.0128	0.0169
AR(2)				0.173	0.164	0.476

Notes. Standard errors in parentheses as indicated in the top column; *** $p<0.01$, ** $p<0.05$, * $p<0.1$. Instrument restrictions are: res=lags restricted to 3 & 4, col=collapsed, ort=orthogonalized. For the specifications in columns 1-3, the misspecfication tests have not been computed because none of the variables of interest are significant. Results for further alternative instrument restrictions are similar (available upon request). Also note that the results in columns 4-6 do not rely on multiple imputation estimation due to problems with varying omitted terms in the interactions. Instead, the data are averaged across the 100 imputations before instead of after the estimation. While this may affect the resulting point estimate of the coefficient, it is highly unlikely that it is qualitatively different. The "true" standard errors are also smaller, which does not, however, challenge the finding that this effect - significant or not - vanishes once separate transition-year effects are added to the model.

Table 2.A.4: GMM results, quadratic level specification

Dep. var.: GDP growth	(1) res	(2) col	(3) colres	(4) ortres	(5) ortcol	(6) ortcolres
Gini(t-1)	0.000389	-0.00695	0.00943	0.000494	-0.00501	0.00859
	(0.00483)	(0.00717)	(0.0275)	(0.00372)	(0.00686)	(0.0226)
Gini(t-1)^2	-3.92E-06	7.09E-05	-8.83E-05	-3.17E-06	5.43E-05	-5.97E-05
	(5.09e-05)	(7.79e-05)	(0.000289)	(4.25e-05)	(7.67e-05)	(0.000252)
Transition*Gini(t-1)	0.00785	0.0281	0.0291	0.00985	0.0195	0.0265
	(0.0170)	(0.0196)	(0.0452)	(0.0179)	(0.0237)	(0.0373)
Transition*Gini(t-1)^2	-6.49E-05	-0.000308	-0.000473	-8.32E-05	-0.000185	-0.000522
	(0.000268)	(0.000304)	(0.000603)	(0.000289)	(0.000371)	(0.000575)
GDP(t-1)	-0.0922***	-0.0760***	-0.187**	-0.0770***	-0.0855***	-0.188**
	(0.0199)	(0.0208)	(0.0763)	(0.0169)	(0.0200)	(0.0841)
PI(t-1)	-0.0117	-0.0138	-0.0112	-0.00691	-0.0134	-0.0527
	(0.0177)	(0.0174)	(0.0609)	(0.0114)	(0.0196)	(0.0566)
Schooling_m(t-1)	-0.0131	-0.024	0.0594	0.00379	-0.0174	-0.00192
	(0.0320)	(0.0275)	(0.0876)	(0.0187)	(0.0253)	(0.0688)
Schooling_f(t-1)	0.0178	0.0138	-0.00509	0.00335	0.0161	0.0297
	(0.0242)	(0.0267)	(0.0544)	(0.0168)	(0.0256)	(0.0482)
F-test of quadratic terms (p-value)	0.4456	0.1237	0.7727	0.2439	0.2759	0.673
Observations	566	566	566	590	590	590
# of countries	115	115	115	116	116	116
Year FE	YES	YES	YES	YES	YES	YES

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1. The second row indicates the type of instrument restriction imposed on the GMM (res=lags restricted to 3&4, col=collapsed, ort=orthogonalized).

Table 2.A.5: Basic specification in differences

Dep. var.: GDP growth	(1) differences	(2) differences with transition country dummies
Δ Gini(t-1)	-0.000206	-0.000111
	(-0.000214)	(-0.000206)
Transition*Δ Gini(t-1)		-0.00069
		(-0.000836)
GDP(t-1)	-0.0635***	-0.0639***
	(-0.0106)	(-0.0106)
PI(t-1)	-0.0140**	-0.0141**
	(-0.00625)	(-0.00635)
Schooling_m(t-1)	-0.00774	-0.00756
	(-0.0105)	(-0.0104)
Schooling_f(t-1)	0.013	0.0127
	(-0.0115)	(-0.0115)
Constant	0.553***	0.556***
	(-0.0866)	(-0.087)
Observations	577	577
# of countries	115	115
Year FE	YES	YES

Notes. Standard errors in parentheses as indicated in the top column; *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 2.A.6: Splines with level Gini, FE results

	(1)	(2)	(3)	(4)
	"change" definition brackets (in percent):			
Dep. var.: GDP growth	3	5	10	20
GDP(t-1)	-0.0505***	-0.0505***	-0.0501***	-0.0495***
	(0.00958)	(0.00967)	(0.00977)	(0.00986)
PI(t-1)	-0.00805	-0.00812	-0.00806	-0.00775
	(0.00897)	(0.00901)	(0.00911)	(0.00878)
Schooling_m(t-1)	-0.00332	-0.00297	-0.00274	-0.00226
	(0.00851)	(0.00850)	(0.00844)	(0.00834)
Schooling_f(t-1)	0.00635	0.00608	0.00575	0.00514
	(0.00945)	(0.00951)	(0.00954)	(0.00943)
Gini_net(t-1)	**-0.000158**	**-0.000137**	**-8.68E-05**	**-2.66E-05**
	(0.000322)	(0.000321)	(0.000320)	(0.000330)
Negative change	**-0.000278**	**-0.000171**	**1.05E-05**	**0.000495**
	(0.000302)	(0.000343)	(0.000469)	(0.000938)
No change	**0.00100***	**0.000358**	**1.59E-05**	**-0.000132**
	(0.000508)	(0.000340)	(0.000211)	(0.000158)
Positive change	**-0.00103*****	**-0.00105*****	**-0.00117*****	**-0.00142*****
	(0.000227)	(0.000238)	(0.000264)	(0.000326)
Constant	0.436***	0.436***	0.432***	0.436***
	(0.0736)	(0.0745)	(0.0762)	(0.0808)
Observations	614	614	614	614
# of countries	115	115	115	115
Year FE	YES	YES	YES	YES

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1. Numbers in the third row represent the knots for defining the "no change"-bracket, i.e., changes between +/-3(5, 10, 20) percent are coded as "no change", and changes above (below) as increases (decreases).

Table 2.A.7: Splines, FE results with sample excluding transition countries

	(1)	(2)	(3)	(4)
	"change" definition brackets (in percent):			
Dep. var.: GDP growth	3	5	10	20
GDP(t-1)	-0.0466***	-0.0468***	-0.0470***	-0.0470***
	(0.0104)	(0.0104)	(0.0105)	(0.0104)
PI(t-1)	-0.00893	-0.00903	-0.00913	-0.00908
	(0.00964)	(0.00959)	(0.00959)	(0.00961)
Schooling_m(t-1)	-0.00948	-0.00979	-0.0101	-0.0102
	(0.00802)	(0.00801)	(0.00806)	(0.00785)
Schooling_f(t-1)	0.0093	0.00974	0.0101	0.0103
	(0.0105)	(0.0105)	(0.0106)	(0.0103)
Negative change	-0.000341	-0.000238	-0.000121	0.000238
	(0.000272)	(0.000313)	(0.000433)	(0.000806)
No change	0.000154	-0.000134	-0.0002	-0.000188
	(0.000452)	(0.000306)	(0.000188)	(0.000134)
Positive change	-0.000154	-0.000129	-0.000119	-0.00017
	(0.000219)	(0.000237)	(0.000291)	(0.000410)
Constant	0.399***	0.402***	0.405***	0.413***
	(0.0842)	(0.0848)	(0.0864)	(0.0888)
Observations	549	549	549	549
# of countries	100	100	100	100
Year FE	YES	YES	YES	YES

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1

Table 2.A.8: 10-year averages, levels and differences

	(1)	(2)	(3)	(4)
Dep. var.: GDP growth	levels & differences	levels & differences with transition countries	levels & differences with transition-countries & trans.-year effects	levels & differences with transition-year effects
Gini(t-1)	0.000309	0.000369	0.000336	0.000306
	(0.000422)	(0.000471)	(0.000472)	(0.000461)
Transition*Gini(t-1)		0.000224	-0.00335	
		(0.00194)	(0.00258)	
ΔGini(t-1)	-2.26E-05	-0.000179	-0.000152	-0.000102
	(0.000253)	(0.000294)	(0.000302)	(0.000282)
Transition*ΔGini(t-1)		0.000702	0.00186	
		(0.000679)	(0.00119)	
GDP(t-1)	-0.0181**	-0.0191**	-0.0153*	-0.0145*
	(0.00891)	(0.00887)	(0.00901)	(0.00859)
PI(t-1)	-0.0397***	-0.0391***	-0.0361**	-0.0358**
	(0.0144)	(0.0145)	(0.0150)	(0.0149)
Schooling_m(t-1)	-0.00511	-0.00496	-0.00602	-0.00568
	(0.00786)	(0.00886)	(0.00927)	(0.00912)
Schooling_f(t-1)	0.003	0.00345	0.00452	0.00399
	(0.00875)	(0.00978)	(0.0103)	(0.0100)
Constant	0.185**	0.189**	0.160**	0.147**
	(0.0756)	(0.0751)	(0.0755)	(0.0726)
Observations	183	183	183	183
# of countries	91	91	91	91
Year FE	YES	YES	YES	YES
Trans-Year FE	NO	NO	YES	YES

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.9: Two-year lag FE and alternative GMM specifications

Dep. var.: GDP growth	(1) FE lag2	(2) diffGMM coll	(3) diffGMM coll	(4) sysGMM res	(5) sysGMM coll	(6) sysGMM coll
Gini(t-1)	0.000686**	1.69E-05	0.000892	0.000102	-0.000485	-0.000171
	(0.000267)	(0.000785)	(0.000785)	(0.000482)	(0.000474)	(0.000346)
GDP(t-1)	-0.0489***	-0.1000***	-0.0633***	0.00777	-0.00127	-0.0017
	(0.00998)	(0.0159)	(0.0147)	(0.00547)	(0.00445)	(0.00344)
PI(t-1)	-0.0142**	-0.0344***	-0.0306**	-0.0181	-0.00928	-0.00663
	(0.00565)	(0.0126)	(0.0123)	(0.0136)	(0.00819)	(0.00772)
Schooling_m(t-1)	-0.00597	-0.0535**	-0.0166	-0.0108	-0.00672	-0.00246
	(0.00849)	(0.0225)	(0.0195)	(0.0110)	(0.00726)	(0.00971)
Schooling_f(t-1)	0.00737	0.0410*	0.015	0.00824	0.00663	0.00339
	(0.0108)	(0.0214)	(0.0222)	(0.0109)	(0.00728)	(0.00942)
Constant	0.412***			-0.0106	0.075	0.0602*
	(0.0788)			(0.0503)	(0.0503)	(0.0337)
Observations	625	481	506	625	625	625
R^2	0.438					
# of countries	119	113	114	119	119	119
Year FE	YES	YES	YES	YES	YES	YES
Trans-Year FE	YES	YES	YES	YES	YES	YES
# instruments		89	97	69	103	111
Hansen Test		0.417	0.136	0.145	0.147	0.223
Sargan Test		1.27E-05	1.39E-06	0	0	0
AR(1)		0.0207	0.00492	0.014	0.00845	0.00751
AR(2)		0.925	0.25	0.201	0.324	0.292

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.10: WIID Ginis (adjusted), Version 2, FE results

Dep. var.: GDP growth	(1)	(2)	(3)	(4)
Gini(t-1)	0.000813**	0.000456	0.0005	0.000368
	(0.000370)	(0.000399)	(0.000386)	(0.000365)
Transition*Gini(t-1)		0.00197	-0.00126	
		(0.00126)	(0.00105)	
GDP(t-1)	-0.0512***	-0.0500***	-0.0410***	-0.0411***
	(0.0103)	(0.0104)	(0.0109)	(0.0110)
PI(t-1)	-0.00801	-0.00783	-0.00952	-0.00951
	(0.00578)	(0.00604)	(0.00730)	(0.00737)
Schooling_m(t-1)	0.00704	0.00874	0.00333	0.00369
	(0.0109)	(0.0116)	(0.00989)	(0.01000)
Schooling_f(t-1)	0.00144	-0.000856	0.000225	-2.30E-06
	(0.0112)	(0.0118)	(0.0116)	(0.0117)
Constant	0.399***	0.398***	0.342***	0.345***
	(0.0846)	(0.0865)	(0.0873)	(0.0878)
Observations	562	562	562	562
R^2	0.33	0.344	0.485	0.481
# of countries	118	118	118	118
Year FE	YES	YES	YES	YES
Trans-Year FE	NO	NO	YES	YES

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.11: IV with OPEC-year effects

Dep. var.: GDP growth	(1) cereal production	(2) agr. land	(3) agr. machinery
Gini(t-1)	-0.00344	-0.00502	0.00135
	(-0.00279)	(-0.00371)	(-0.00334)
GDP(t-1)	-0.0438***	-0.0402***	-0.0683***
	(-0.0137)	(-0.0151)	(-0.0125)
PI(t-1)	-0.00511	-0.00351	-0.0260**
	(-0.00589)	(-0.00662)	(-0.0124)
Schooling_m(t-1)	-0.00869	-0.0129	0.022
	(-0.0118)	(-0.0138)	(-0.02)
Schooling_f(t-1)	0.00465	0.00587	-0.0122
	(-0.0106)	(-0.0119)	(-0.0187)
Cereal production*instr	3.59E-09		
	(2.05E-09)		
Cereal production	-2.35E-09		
	(-1.70E-09)		
Agr. land *instr		8.85E-06	
		(-7.74E-06)	
Agr. land		-3.81E-05	
		(-0.000798)	
Agr. machinery*instr			-2.17E-07
			(-1.89E-07)
Agr. machinery			3.04E-06
			(-4.24E-06)
Observations	584	584	346
# of countries	92	92	75
Year FE	YES	YES	YES
Transition-Year FE	NO	NO	NO

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1.

Table 2.A.12: IV with transition countries

Dep. var.: GDP growth	(1) basic	(2) trans-year	(3) transinteract	(4) transinteract & trans-year
Gini(t-1)	-0.00222	-0.000594	-0.000558	-0.000886
	(-0.00152)	(-0.00148)	(-0.00176)	(-0.00162)
Transition*Gini(t-1)			-0.0073	0.00417
			(-0.00593)	(-0.0035)
GDP(t-1)	-0.0489***	-0.0413***	-0.0652***	-0.0378***
	(-0.0127)	(-0.0102)	(-0.0188)	(-0.0115)
PI(t-1)	-0.00363	-0.00588	-0.00578	-0.00548
	(-0.00603)	(-0.00567)	(-0.00579)	(-0.00584)
Schooling_m(t-1)	-0.00604	-0.00237	-0.0106	-0.00222
	(-0.00808)	(-0.00614)	(-0.00972)	(-0.00621)
Schooling_f(t-1)	0.00469	-0.00137	0.014	-0.00221
	(-0.00842)	(-0.00661)	(-0.0116)	(-0.00684)
Observations	566	566	566	566
# of countries	92	92	92	92
Year FE	YES	YES	YES	YES
Transition-Year FE	NO	YES	NO	YES
OPEC-Year FE	YES	YES	YES	YES
FIRST STAGE				
GDP(t-1)	3.196**	4.286***	3.165**	4.331***
	(-1.385)	(-1.323)	(-1.392)	(-1.324)
PI(t-1)	0.7	0.752	0.71	0.783
	(-0.705)	(-0.69)	(-0.706)	(-0.692)
Schooling_m(t-1)	-2.258*	-1.603	-2.207	-1.594
	(-1.366)	(-1.376)	(-1.347)	(-1.377)
Schooling_f(t-1)	0.962	-0.0476	0.922	-0.0733
	(-1.537)	(-1.544)	(-1.521)	(-1.546)
SWratio*instr(t-2)	0.294***	0.289***	0.281***	0.278***
	(-0.0856)	(-0.0931)	(-0.0948)	(-0.0951)
Transition*instr(t-2)			0.0517	0.369
			(-0.154)	(-0.277)
Observations	566	566	566	566
R^2	0.16	0.228	0.16	0.228
# of countries	92	92	92	92
Weak instruments: F-stat	11.8	9.64		
Kleibergen-Paap max. bias			15%	15%

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1.

Chapter 3

An Index of Inter-Industry Wage Inequality: Trends, Comparisons, and Robustness

3.1 Introduction

This paper introduces a newly constructed measure of wage inequality. A Theil index is computed for manufacturing sectors across a large number of countries and a time period of up to 48 years. The index itself is very similar to the one developed in Galbraith et al. (1999) and Conceição and Galbraith (2000). As part of the University of Texas Inequality Project (UTIP), they constructed a Theil index based on the same data source employed here, resulting in the UTIP-UNIDO measure of wage inequality. Building on the work of the UTIP, several concerns regarding the treatment of the raw data as well as questions of internal and external validity which have remained open up to this point are addressed in this paper on the basis of the newly constructed index.

The index is based on sectoral data from the UNIDO Industrial Statistics at the 2-digit level of International Standard Industrial Classification (ISIC, Rev 3), which is a fairly crude level of aggregation. One of the main reasons why this dataset is attractive for the construction of an inequality measure is its broad time and country coverage. If the resulting, narrowly defined index of wage inequality is able to mirror overall changes in income inequality, as argued by Galbraith and Kum (2005), it can be applied in many additional contexts than just for analyses of wage inequality (or manufacturing wage inequality, for that matter). Importantly, it could serve as a proxy for developments in overall income inequality in empirical applications focusing on changes over time, such as the fixed effects model typically employed in country-level macro panel regressions. This paper tests whether the broad generalizability and applicability claimed for the UTIP index also holds for the Theil index constructed here, which is shown to be fairly similar to the UTIP-UNIDO index for many of the countries covered by both measures. A detailed comparison of the new index with the one developed by the UTIP is provided and I try

Acknowledgements: I would like to thank Stephan Klasen, Axel Dreher, and Sebastian Schneider for helpful comments and suggestions on this essay, and Anne Körner for excellent research assistance.

to explain observed differences between the two measures. Building on prior work by Galbraith and Kum (2005), and Galbraith et al. (2015), the "external" validity of the index, that is, the extent to which the Theil index is representative of overall income inequality, is then examined. The new index is subjected to a number of comparisons with measures of income inequality to check whether it can predict developments in overall inequality. Unfortunately, the results do not lend much support to this idea. The estimates suggest that the association between the Theil index of manufacturing and income inequality is neither very stable, nor strong enough to postulate an economically meaningful link between the two concepts.

Doubts arise not only for the "external validity" of the index, but also concerning its "internal" capability to accurately reflect developments in manufacturing wage inequality. Because the index relies on sector-level data on wages and employment which are aggregated at the 2-digit level of industrial classification, it only measures between-sectoral wage inequality and cannot give an account of inequality within sectors. Since the Theil index is decomposable into inequality between and within units, the availability of less aggregated data for subsectors at the 3- and 4-digit level allows the calculation of part of the within-component for a smaller subsample of countries and years. Despite the fact that inequality within the most detailed sectoral classification available still remains unaccounted for - which, given an average number of almost 22,000 employees in the smallest unit of record, is likely to be substantial - I find within-sectoral inequality to make up at least 40% of overall manufacturing wage inequality. While it is obvious that manufacturing wage inequality at a single point in time is vastly underestimated, (Conceição and Galbraith, 2000) argue that the between-sectoral index is still able to trace changes in within-sectoral inequality over time. However, despite the unbalancedness of the 3- and 4-digit level data, I find indication that in around 13% of cases, a Theil index relying on the between-sectoral component of manufacturing wages conveys an incorrect image of overall changes in manufacturing wage inequality. Given that the "true" extent of within-sectoral inequality is likely to be considerably larger because inequality between individuals within subsectors still remains unaccounted for in the more detailed data, this number provides a lower bound to the true discrepancy between between-sectoral and overall changes in manufacturing wage inequality.

Before moving into the analysis of the final index, the paper provides a thorough description of the challenges inherent in the raw data for creating a consistent measure of inequality over time, and the strategies employed to deal with them. The main problem with exploiting the UNIDO industrial statistics for inequality measurement is the unbalancedness of the sectoral data. For computing an index of wage inequality, data on employees and wages are used, and in order to obtain meaningful and comparable values over time, both variables have to be present in every year in all of the sectors included in the measure. To arrive at any useful measure of inequality, some modifications of the raw data are therefore unavoidable. In the initial documentation of the UTIP-UNIDO index, it is not apparent how this was handled in the construction of the measure. In the meantime, documentation has improved slightly with the recent release of the accompanying paper (Galbraith et al. 2014) to the update of the index in 2013. But, as the authors themselves state, "These issues were handled on a case-by-case basis, using judgement and common sense to arrive at a set of "final revised values"." (Galbraith et al. 2014: 2). I received the information from the authors that, apart from differences in the imputation of missing values, deviations between my newly computed index and the most recent version of the UTIP measure stem from the fact that the UTIP has harmonized their index to match up with its previous versions (no longer available on the UTIP webpage). However, it is not clear how large the adjustment has been in specific cases, and which countries and/or years are affected. In a direct comparison of the index computed in this paper with that provided by the UTIP, I find several countries covered by their index for which the version of the UNIDO industrial statistics which this paper is based on does not provide data. It is therefore clear that the inequality series for entire countries is based on a different version of the data.

Ultimately, for achieving balancedness of the sectoral data, the choice lies between dropping sectors with poor data coverage or the imputation of the missing values. I use both strategies in this paper, and which one I choose for a particular case depends on the number of missings as well as the number of years which would be lost due to the inclusion of a sector with limited time coverage. Testing the robustness of the measure to the discretionary - and somewhat arbitrary - decision of whether and when to impute a value, and/or to drop a sector, is not straightforward. While the UTIP data are an obvious benchmark, I also employ a few other strategies of testing the index' robustness "internally," as described in section 3.2.3 below. The imputation methods used are described in section 3.2.4, and appendix 3.B contains a more detailed version with examples from the raw data to illustrate some of the cases typically encountered. As a rule of thumb, no sector is included in the final measure which contains more than 50% imputed data points. The imputation of data points on the raw data is regarded as relatively unproblematic since it relies on "internal" data sources originating from the UNIDO industrial statistics only. Furthermore, because imputation is performed at the sectoral level, the impact of each single imputation on the final index is very small given that the UNIDO data include up to 23 manufacturing sectors per country. While there is no way of knowing exactly the impact of imputing on the final measure, a variable is retained which contains the number of imputed data points in the underlying sectors in every year. Including it into the comparison of the Theil index to other inequality measures at least enables a judgement of whether or not it makes a difference if the index relies on imputed data.

The effect of the dropping of sectors on the accuracy of the resulting "long" inequality measure, both in terms of inequality levels and changes over time, can be assessed more easily. For those years in which data for dropped sectors are available, the Theil index is computed with and without these sectors and the resulting values are compared. As shown in section 3.2.3, the impact of the dropped sectors is very limited in most cases. Whenever the deviation between the two numbers is larger than 10%, an alternative version of the Theil index is computed which comprises more sectors and therefore provides more accurate numbers for inequality levels. However, the resulting "short" index covers fewer years and therefore compromises the original advantage of the index, which was to provide an account of the developments in wage inequality over long periods of time. Therefore, whenever the "long" version accurately traces developments over time, despite providing unreliable numbers for inequality levels, a recommendation is made to retain the version for dynamic analyses of inequality - that is, empirical applications which focus on changes of inequality over time. This results in a preferred version of the index with an average time coverage of 28.5 years, which I label the "dynamic" version and on which the remainder of the paper is based.

Section 3.2 is dedicated to the construction of the index and the associated preparation and treatment of the raw data: the index and its mathematical properties are introduced in part 3.2.1. Part 3.2.2 is concerned with the sectors used for the construction of the index. General information and descriptives are provided in part 3.2.3, along with information on the properties of the dropped and retained sectors, and an analysis of the effect of dropping sectors on the accuracy of the final measure. Part 3.2.4 deals with the sectors entering the measure and describes the imputation methods used for attaining balanced data. Section 3.3 provides some basic descriptives and information on the index in part 3.3.1 and makes a comparison to the UTIP index in part 3.3.2. Section 3.4 focuses on the role of within-sectoral inequality and section 3.5 relates the index to measures of income inequality. Section 3.6 concludes.

3.2 Constructing a Theil Index of Inter-Industry Wage Inequality

3.2.1 The Theil index

To compute the Theil index of inter-industry wage inequality, I make use of the data on total wages and employment that are provided at the sector level for a maximum of 23 manufacturing sectors, as per the ISIC 2-digit sectoral classification. The between-sector component of the Theil is defined as

$$T' = \sum_{s=1}^{S} y_s \cdot ln(\frac{y_s}{n_s})$$

with S denoting the different sectors, s=1, ..., S. y_s represents a sector's wage share, defined as a sector's wage bill divided by the sum of wages of all industries, while n_s represents the "population" (=employment) share of sector s, defined as the sector's employment over total employees (Theil 1967). This original representation of the index in shares[1] is not as common, yet it is insightful because it makes it easy to illustrate several properties of the index.[2] First, the sector's wage share can be interpreted as the weight with which each sector enters the measure. Second, if the ratio of the wage share and the population share are equal, taking their logarithm yields zero, which implies that the sector does not enter the measure. Consequently, if all income shares and population shares are equal, the between-group Theil takes its lower bound value of zero, indicating

[1] As opposed to the representation in averages, which is mathematically equivalent.

[2] For a more detailed discussion of the properties of the Theil index, consult Conceição and Ferreira (2000).

a perfectly equal distribution of income between sectors. Although the contribution of a sector to the measure will be negative whenever the population share is larger than the income share, multiplying the log value with the income share ensures that positive values have a larger weight in the final measure. This is because for every unit that has a smaller income- than population share, there must be at least one for which the opposite is true. Because positive values by construction result from income shares larger than population shares, the positive values will automatically be multiplied with a larger number than the negative ones. T' can therefore never be negative. The measure has no upper bound, which makes intuitive interpretation of a single number difficult, but comparing numbers based on the same underlying units - in this case, industrial sectors - is straightforward. Although the index is sensitive to the numer of underlying sectors S, it can be easily normalized by dividing the value by its theoretical maximum, log S. A variant of the generalized entropy class of indequality measures, the index furthermore has the adavantage of being perfectly decomposable into an infinite number of fractals, each representing within-unit inequality at a lower (i.e., more disaggregate) level. Since UNIDO also provides data at more detailed levels of sectoral aggregation (3- and 4-digit level) for some years, the use of the Theil index enables a judgement of at least part of the extent of within-sectoral inequality as compared to between-sectoral inequality. The formula including within-sectoral inequality at the 3-digit and 4-digit levels is as follows:

$$T' = \sum_{s=1}^{S} y_s \cdot (ln(\frac{y_s}{n_s}) + \sum_{s3d=1}^{S3d} y_s 3d \cdot (ln(\frac{y_{s3d}}{n_{s3d}}) + \sum_{s4d=1}^{S4d} y_s 4d \cdot ln(\frac{y_{s4d}}{n_{s4d}})))$$

y_{s3d} represents the share of each 3-digit sector's wage in their respective 2-digit sector, and y_{s4d} is the share of each 4-digit sector's wage in their respective 3-digit sector. Equivalently, n_{s3d} and n_{s4d} are the corresponding employment shares. A detailed discussion of the within-sectoral decomposition and its limitations can be found in section 3.4. Before moving to a discussion and analysis of the between-sectoral component of the index, the next sections describe the procedures used for achieving balanced versions of the underlying raw data, which is a prerequisite for obtaining values of the Theil index that can be meaningfully compared over time.

3.2.2 Between-sectoral inequality

The main challenge in exploiting the UNIDO industrial statistics for inequality measures is unbalancedness, both between sectors and over time. In order to obtain meaningful and comparable values over time, the same sectors should be included in the inequality measure every year in a given country. Hence, if data for one sector is missing in only one out of the 48 years, this means that either that year needs to be dropped, or the sector must be excluded from the index in all of the remaining 47 years. This poses great challenges given the highly unbalanced nature of the raw data. The problem is exacerbated with the inclusion of lags in empirical applications, which is typically done in macroeconomic regressions with inequality as the dependent variable due to the high degree of inertia in the measure. Already a one-year gap leads to the loss of at least

2 data points in the estimation sample, and data for single years (i.e., with missings in both the previous and subsequent year) drop out altogether. ection In order to obtain a workable index which can be readily used in empirical analyses, some imputation as well as interpolation is therefore indispensable. The choice between imputing missing values and dropping sectors is effectively a trade-off between two objectives. On the one hand, one wants to maximize time coverage - in particular, to fill short gaps within longer spells of data. On the other hand, the loss of information arising from the dropping of sectors should be minimized in order to ensure accuracy of the resulting inequality statistics.

It should be mentioned that the assumption is that missings in the underlying data are random across sectors. There are no patterns in the raw data suggesting otherwise, and the fact that in most instances, data is missing only in a few sectors and often in only one dimension - wages or employees - supports this view. Whether this is also true for entire years of missing data is not as clear. Because the UNIDO industrial statistics rely on surveys from establishments, the fact that no data was compiled in a certain year might have reasons which could also affect the economy as a whole, including the manufacturing industry.[3] However, there is no reason to expect that industrial sectors are affected asymmetrically and that inequality in those years in which data is missing is very different from that in the preceding and subsequent years. The following paragraph provides a brief overview of the sectors covered in the UNIDO Industrial Statistics, shows the impact of dropping sectors on the inequality index, and offers a solution on how to treat those cases where large differences arise between indices with and without dropped sectors. Section 3.2.4 will then focus on the retained sectors and describe how missing values have been dealt with.

3.2.3 Dropping sectors

Although there is a trade-off between time- and sectoral coverage - with the former implying a loss of accuracy in the resulting inequality measure's ability to capture between-sectoral wage inequality - it is much less severe than one might initially expect. It turns out that in most instances, those sectors which are not well covered by the data are also the ones which are of lower economic significance for a country, and hence are also relatively small. Because the Theil index weighs the logged discrepancy between wage and employment shares by each sector's wage share, this means that the smaller sectors are also relatively less important in determining the final value of the index. Hence, omitting these sectors often changes the index very little. Before moving to a systematic analysis of the effects of dropping sectors, table 3.1 provides an overview of the 23 manufacturing sectors covered by the data and provides information about their average size (as measured by the wage share), the discrepancy between the wage- and employment shares, and the total number of times each sector has been included and excluded for the "long" version of the index, which aims at maximizing time coverage.[4]

[3] While the documentation of the Industrial Statistics database contains a detailed description of how non-response for individual establishments was dealt with, there is no mentioning of why entire sectors or even years are missing in some countries.

[4] Information on dropped sectors for individual countries can be found in appendix table 3.A.9. All numbers presented rely on the balanced version of the data, i.e., including the imputed data points.

Clearly, the most frequently dropped sectors are 19, 30, 32, 35, and 37. They are available only for later time periods (1990s onwards) and their inclusion would therefore mean a substantial loss in time coverage, especially when the time series is long and covers a lot of the early years.[5] Luckily, these sectors tend to be relatively small on average, with wage shares ranging from 0.2 to a maximum of 2.9 percent of total manufacturing wages.

Apart from the wage share, the second aspect determining the importance of a sector for the Theil index of wage inequality is the discrepancy between the wage- and the employment share. Here, the omitted sectors cover a broad spectrum, with sectors 19 and 37 having a lower wage- than employment share and the rest having 25-37 percent larger wage- than employment shares. These sectors therefore do contribute to inequality, but because each contribution is weighted with a relatively small wage share, their final contribution will be rather small.

Of course, there are other omitted sectors, and one might worry, for example, about the exclusion of sector 23 in 15 cases, which is the sector with the highest average discrepancy between the wage- and employment share, and has a wage share of 3.4 percent. Furthermore, the low average size of the frequently omitted sectors does not mean that this is also the case in an individual country, and some sectors might be of high economic significance in single economies.

As a general check of the degree to which the "long" version of the index, wherein sectoral coverage has been sacrificed for the sake of a longer time series, is representative of the overall level and development of between-sectoral wage inequality, I have therefore also computed the index for every country and year using all of the available data, including those for the dropped sectors. The resulting "full" index is not comparable over time, but it can serve as a benchmark for the comparison with the long version. The percentage difference between the two measures serves as a first indication of the degree of distortion introduced by the omission of certain sectors. Averaging over all the available countries and years,[6] the two versions seem rather similar, with the "long" version yielding 10.6% higher inequality numbers on average across all countries and years. This rather low average deviation[7] is, however, concealing large variations across, as well as within, countries. While the two versions are virtually identical in a large number of countries, others display a large difference between the indices. Moreover, in a substantial number of countries which have a low average difference, there is a lot of variation over the years. Appendix table 3.A.1 displays the deviation between the two versions of the Theil index for all countries where the indices differ, sorted by the maximum percentage deviation.[8] In addition to the maximum, the table also reports the mean percentage deviation, and the

[5]The reason for this is the change of the ISIC classification scheme from Rev. 2 to Rev. 3 in 1989, and the accompanying re-categorization of old industries, and creation of new industrial categories such as, e.g., sector 37 (Recycling).

[6]Only years when deviations actually occur between the two indices have been included in the computation of the different measures of convergence. In the dataset, a deviation of 0 arises if, and only if, the sectoral coverage is the same between the measure and including those years would skew the similarity indicators upwards.

[7]While this number may not appear as very small at first glance, it is driven upwards by a few "outlier" countries with very high mean deviations of above 100%

[8]The mean deviation is based on the absolute value of the negative deviations, i.e., cases in which the long version is larger than the full one.

Table 3.1: Overview of manufacturing sectors

ISIC code	Manufacturing sector (ISIC Rev. 3)	Number of times... included	excluded	total	Wage share	Average ratio of wage- to employment share	Technology level
15	Food products and beverages	113	0	113	22.7	102	Low
16	Tobacco products	96	9	105	2.4	176	Low
17	Textiles	111	2	113	10.7	85	Low
18	Wearing apparel; dressing and dyeing of fur	105	5	110	8.3	73	Low
19	**Tanning and dressing of leather; luggage, handbags, saddlery, harness & footwear**	**12**	**62**	**74**	**23**	**72**	Low
20	Wood and of products of wood and cork, excl. furniture; articles of straw and plaiting materials	109	3	112	36	80	Low
21	Paper and paper products	109	3	112	25	118	Low
22	Publishing, printing and reproduction of recorded media	103	9	112	39	124	Low
23	Coke, refined petroleum products and nuclear fuel	89	15	104	34	271	Medium-Low
24	Chemicals and chemical products	109	3	112	76	144	Medium-High
25	Rubber and plastics products	101	8	109	35	104	Medium-Low
26	Other non-metallic mineral products	109	3	112	64	110	Medium-Low
27	Basic metals	102	8	110	44	149	Medium-Low
28	Fabricated metal products, except machinery and equipment	106	7	113	52	103	Medium-Low
29	Machinery and equipment not elsewhere classified	100	9	109	42	110	Medium-High
30	**Office, accounting and computing machinery**	**9**	**52**	**61**	**11**	**137**	Medium-High
31	Electrical machinery and apparatus not elsewhere classified	98	10	108	38	116	Medium-High
32	**Radio, television and communication equipment and apparatus**	**9**	**54**	**63**	**29**	**128**	Medium-High
33	Medical, precision and optical instruments, watches and clocks	91	13	104	08	111	Medium-High
34	Motor vehicles, trailers and semi-trailers	96	11	107	41	122	Medium-High
35	**Other transport equipment**	**10**	**57**	**67**	**27**	**125**	Medium-High
36	Furniture; manufacturing not elsewhere classified	107	5	112	36	83	Low
37	**Recycling**	**8**	**50**	**58**	**02**	**84**	Low

standard deviation. These figures can provide a broad idea of the "static" resemblance of the long- with the full version of the Theil index. Researchers who care less about the level of wage inequality, but are rather interested in its development over time - which is arguably one of the main advantages of this dataset - may be more concerned about the ability of the index to trace changes in inequality. Table 3.A.1 therefore also includes the correlation of both the level and the differences of the "long"- with the "full" index.[9] Apart from Kuwait, which has a correlation of 0.89 in differences, none of the countries with a level deviation of less than 10% has a correlation lower than 90%, neither in levels nor

[9]Only looking at the correlations gives a slightly more optimistic, but qualitatively similar picture. Those countries showing lower percentage deviations of inequality levels between the long and the full version of the index generally have higher correlations as well, but not necessarily vice versa. Senegal, for example, has a correlation of 0.998 between the two indices over the 28 years.

differences.[10] The same applies to the mean deviation, which - apart from the Philippines - is always below 6.2 percent. Interestingly, in several cases where the deviation of the level of the two indices is rather large (e.g., Botswana and the Netherlands), the correlation between the two indices is still high (above 0.92). Despite starting from very different levels of wage inequality, changes over time seems to still be well captured by the long version of the index in several cases.[11]

In an attempt to address these issues and provide more accurate versions of the index, I therefore recalculated the index including more of the previously omitted sectors, with the same constraint of including the same sectors in all years. Naturally, this implies losing several years of data given that the initial motive behind excluding the sectors was to increase time coverage. In many cases, this leads to the inclusion of all sectors, but there are still sectors which are excluded also from these "short" versions. I have calculated short versions for all countries with a deviation of more than 10% in any year (as indicated by the maximum)[12] and then repeated the above exercise (results are shown in panel 2 of table 3.A.1).[13]

Of course, nothing can be said about the counterfactual deviation in the years for which data is missing on the sectors which have been dropped from the index, but there is no reason to suspect that the deviation would be larger than in the years covered. Specifically, I checked whether there is a discernible time trend in the deviation over the years, and it does not seem to be the case that the contributions of omitted sectors is growing or decreasing over time. The contributions are also not varying in any other systematic manner which would allow inference about their development outside the sample range. Overall, given the possibility to combine the two versions of the index provided for countries with a deviation of more than 10% between the long and the full version, the Theil index is able to provide an accurate picture of the extent of between-sectoral wage inequality in manufacturing. It then depends on the purpose of the research which version is preferable: those applications of the index for which the development over time is of interest may still benefit from the long version despite larger differences in the levels, and vice versa. The last column of table 3.A.1 provides a recommendation of which index to use in dynamic applications. Since the main purpose of constructing the Theil index was its ability to trace inequality changes over a longer time horizon, I decided to keep this

[10] The correlation is based on only those years with non-zero deviations in order to not artificially drive the correlations upwards.

[11] While the reverse case can also be true, there are only two countries - Brazil and Algeria - which have a high similarity of the inequality levels (less than 10% deviation on average), but a low correlation of the indices over time.

[12] Apart from it being the strictest criterion, I focus on the maximum percentage deviation for another reason: Because not all omitted sectors are always present at the same time, if there is a deviation between the two indices, this is not necessarily the "full" deviation. For example, of, say, 5 sectors which are not covered by the "long" index, only 1 might be included in a given year in the "full" index. If, for that reason, the deviation is lower in years where fewer sectors are present in the "full" version as well, taking the maximum deviation will provide a more accurate indication of the potential bias arising from the omittance of sectors.

[13] In countries with remaining deviations, i.e., where some sectors are still being dropped in the short version, the differences between the full- and the short version are now well below the 10% cut-off. There are a few exceptions where the long version is retained despite larger deviations. They are marked with an asterisk in appendix table 3.A.1 and the reasons for keeping them are explained in detail for every case below the table.

"dynamic" version of the index as the preferred version for the remainder of the paper. In particular, part 3.3.1, describing the final index, and section 3.5 which analyzes its similarity to other inequality indices, are based on the "dynamic" version of the index.[14]

3.2.4 Retaining sectors: Imputation

Even after dropping sectors with low data coverage, the remaining dataset is far from balanced. There are a lot of observations where only one of the two variables, wages and employees, necessary for the index is provided. In other years, both variables are missing in certain sectors. The remaining missings are therefore imputed in order to attain a workable inequality index. It should be noted that due to the extremely heterogeneous data coverage across variables, countries, and years, it is impossible to apply the same imputation procedure to all countries, let alone sectors. There are different ways to impute missing values, with varying degrees of sophistication, and which one is most suitable has to be decided on a case-by-case basis.

The preferred method here is a regression-based approach. I prefer this approach over other imputation methods because it allows exploiting other information from the UNIDO industrial statistics to predict a missing value. Especially in years where there are large changes in wage- or employment shares, simply interpolating values without consulting other information provided in the dataset may lead to suboptimal outcomes and erratic movements in inequality numbers due to large changes in relative sector shares.[15] Two more variables, output and the number of establishments, are provided at the sectoral level and are positively associated with both the number of employees and their (total) wages in a given sector. Their development can be indicative of changes in those variables for which information is missing, and indeed, the relationship between these variables is very strong in many instances. Additionally, often only one of the two variables needed for the computation of the index is missing. In these cases the other one is used in the regression as well (e.g., if a value exists for employees but not for wages, the "employees" variable enters as one of the predictors of wages). Finally, a time trend in the development of wages or employee numbers is sometimes discernible and is also included in the set of potential regressors. The fitted value from a simple OLS of the following exemplary form is used to fill the missing value (in this case for wages):

$$\mathrm{Wages_t} = \alpha + \rho \mathrm{Employees_t} + \beta \mathrm{Establishments_t} + \gamma \mathrm{Output_t} + \delta_\mathrm{t} + \epsilon_\mathrm{t}$$

Again, the main obstacle to the use of this more sophisticated imputation method is data availability. It is not possible to always use the same regressors across countries or sectors, with available variables differing even within the same sector between years. The above example therefore only represents the most general specification while many of the

[14] Because the "dynamic" version ensures accuracy in capturing changes over time and the comparison with other inequality measures is based on fixed effects models, using the "dynamic" version of the index is considered as unproblematic.

[15] For example, if a sector's employment numbers drop drastically in one year and the information on wages is missing, simply linearly interpolating the value for wages based on the previous and next year's value would lead to a large change in the relative ratio of the sector shares, whereas taking into account the information on employment and adjusting the wage value downwards leads to a smoother series.

actual regressions only contain a subset of the variables.

Sometimes there is no further information available at all for a missing observation, or predicting fitted values is not feasible for other reasons (e.g., due to a too-short time period which leaves no degrees of freedom for estimation). In this case, alternative imputation methods have to be explored. Second-best solutions employed in this paper are a simplified hot-deck type approach,[16] where an observation similar to the missing is used, or linear interpolation based on the surrounding values. All methods are described in detail in appendix 3.B, starting with the regression approach.

3.3 Inter-Industry Wage Inequality: Trends and Comparisons

3.3.1 Trends in between-sectoral wage inequality

Before moving to a comparison of the newly computed index with that constructed by the UTIP, a few facts and figures of the index constructed and discussed so far are presented below. All graphs and figures are based on the "dynamic" version of the index as developed in section 3.2.3, which is based on the balanced sectoral dataset after imputation. The overall development of wage inequality is depicted in figure 3.1 below. The first graph shows the overall evolution of the index over a time period of almost 50 years and the second graph breaks it down into developing and developed countries (as per the World Bank GNI threshold definition). Note that in order to make inequality numbers comparable between countries, which differ in the number of sectors underlying the measure, the graphs rely on the normalized version of the Theil index. It is clear that between-sectoral manufacturing wage inequality has been increasing over the sample period, but the largest increase occurs in the second half of the 1980s and the early 1990s. Inequality is higher in developing countries throughout the entire time period and the two series develop rather similarly. This is in line with Galbraith and Kum (2005), who find the same patterns for the first version of the UTIP dataset with data until 1999. Breaking the data down by region, as shown in figure 3.2, is more informative in terms of differential developments across country groups.[17]

It becomes apparent that although a small spike around 1990 appears in several country groups, the large increase from 1980 to 1990 seems to be driven to a large extent by the Middle East and North African (MENA) region (comprising both developed and developing countries). Within this group of countries, it is Tunisia and Kuwait which show very large increases in the late 1980s (shown in appendix figure 3.A.1). The country means of the normalized (Theil(n)) and non-normalized Theil index are compiled in appendix table 3.A.2 along with the main outcomes of the robustness exercises from sections 3.2.3 and 3.2.4, i.e., the number of sectors included the measure in each country and the number of imputed data points. Besides the basic information, which is provided for the preferred,

[16] See Andridge and Little 2010 for a review of the method.

[17] The regional grouping relies on the World Bank classification, but Europe and North America have been pooled together into one category due to the small number of countries in the former.

Figure 3.1: Evolution of the Theil index over the sample period

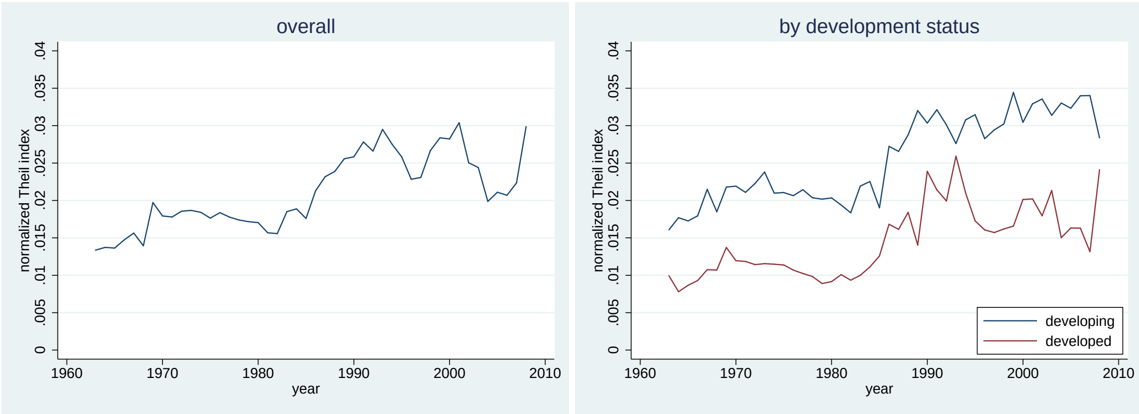

Notes. The first graph is based on a (relatively) constant sample of the 56 countries with a minimum time coverage of 30 years to avoid fluctuations in the time series caused by countries entering or exiting the sample in certain years. Similarly, years with fewer than 30 data points are not shown (affecting the years after 2008). The same years are omitted in the second graph for the same reasons and to ensure comparability of the two graphs.

"dynamic" version for all countries covered by the index, it contains the standard deviation of the (non-normalized) version in the last column to give an idea of the variation of the index within a country. The overall impression from the country averages is in line with the usual country rankings in terms of inequality: the lowest numbers are found in Europe, especially the Scandinavian countries, and high numbers are most prevalent in developing countries and countries from the Middle East. There are a few surprising cases though, such as The Gambia, Nicaragua, and Afghanistan scoring very low on manufacturing wage inequality and Romania, a former communist country, scoring very high. This already provides some indication that manufacturing wage inequality is not always very closely related to a country's overall income inequality, and can sometimes generate a misleading image if such generalizations are drawn. The relationship between manufacturing wage inequality as constructed here and overall income inequality is examined more systematically in section 3.5. Before moving to the question of internal and external validity of the newly constructed index, I compare it to the UTIP index. The more similar the two measures turn out to be, the more will the results from the subsequent validity analysis also apply to the UTIP index. Furthermore, the comparison with the UTIP index can yield some indication as to whether the extent of imputation systematically distorts the resulting measure.

3.3.2 Comparison to the UTIP-index

The first noticeable difference between the newly constructed index and the Theil index calculated by the UTIP is country and time coverage. The average time coverage of the new Theil index is 28.5 years,[18] vs. 26.2 years for the UTIP one, and has information for 137 countries, whereas the UTIP index covers a total of 154 countries. There are 2 countries (Liberia and Serbia) with a total of 15 observations which are in the new index

[18]For the "long" version, average time coverage is 31.3 years.

Figure 3.2: Development of the Theil index by region

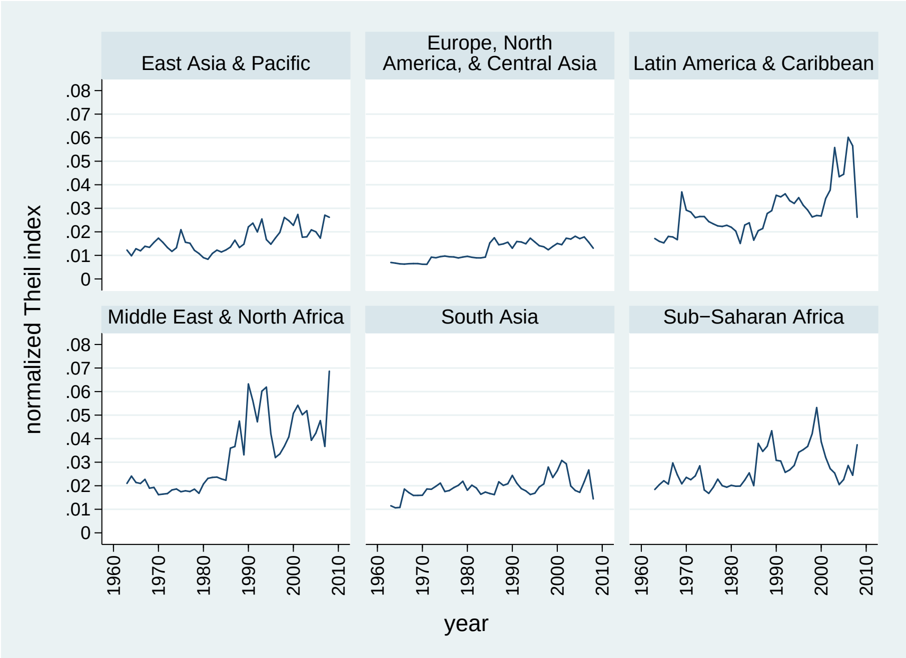

Notes. The graphs are subject to composition effects due to the entering and/or exiting of countries throughout the sample period. The regional subsamples have not been restricted any further to improve time consistency in country coverage because of the lower number of countries per group as compared to the graphs in figure 3.1

but not in the UTIP one, and 19 countries with a total of 210 observations for which the UTIP provides information but which are not covered by the new index. 10 out of these 19 countries are not part of the UNIDO industrial statistics,[19] and another five countries (Armenia, the Bahamas, Rwanda, Sudan, and Zimbabwe) have not been included in the new index due to the lack of useable raw data.[20] This implies that the UTIP indices are based on older versions of the UNIDO data for those countries.[21] Because the older UNIDO data rely on different industrial classification schemes, the index values do not necessarily compare easily from the new to the old versions. In particular, it appears that the previous version of the UTIP index was based on a more detailed, 3-digit level of classification, which makes it more accurate in capturing manufacturing wage inequality. It is not clear from the documentation of the UTIP index in which cases other versions of UNDIO industrial statistics were consulted, and what differences arise from comparing values based on the different industrial classification schemes. It is also not clear when and how the data were harmonized with the previous version of the index. I can therefore not ultimately determine whether differences between the new index and the UTIP one arise due to differential sectoral coverage or varying data sources.

[19] These are Bahrain, Bhutan, Cap Verde, Czechoslovakia, the German Democratic Republic, West Germany, Equatorial Guinea, Myanmar, the Seychelles, and Togo.

[20] That is, although some data is provided for these countries in the UNIDO industrial statistics, the data never covers both wages and employees at the same time.

[21] The other four countries covered by the UTIP but not by the new index have been excluded due to insufficient time coverage. Angola, the United Arab Emirates, Bosnia and Herzegovina, and Cambodia have a maximum time coverage of four years in the UNIDO industrial statistics, of which a maximum of two years are consecutive. The resulting inequality measure would therefore be of little use for comparisons over time, which is the main selling point and the reason for constructing the index in the first place.

The overall correlation between the UTIP and the newly constructed Theil index across all 135 countries covered by both indices is 0.83 for inequality levels and 0.79 for changes in inequality. Although the two indices appear to develop rather similarly on average, the correlations by country reveal large differences and range from a perfect correlation of 1.0 in 17 countries to negative correlations in Bulgaria, Germany, Estonia, Jamaica, and Uganda.[22] Appendix table 3.A.3 displays the number of imputed data points, the year coverage for the UTIP and the "dynamic" version of the Theil index, as well as the correlations and relative deviations of the two measures in levels and differences. The degree of divergence between the two measures is weakly, but significantly correlated with the extent of imputation[23] in the new index. This makes perfect sense given that the reason for imputing values was that the raw data was not utilizable and hence the construction of any sort of index requires a choice of whether to impute or not, and, if applicable, of the imputation method. Obviously, if no imputation is carried out, differences in the measures are implied. But even if the data have been modified in some way, the outcome is not necessarily the same and the resulting indices are likely to still differ to some - smaller - extent. On average, the dynamic version of the Theil index is 3.8 percent higher than the UTIP measure whereas the long version is 1.2 percent lower. While these averages again differ substantially across countries, neither the dynamic, nor the long version display significant correlations between the average sign of the deviation and the number of imputed data points at the country level.

Looking at those countries which display low correlations or very high deviations from the UTIP in more detail, a few peculiarities are noticeable. First, the association with the number of imputations is not stronger in the countries displaying negative correlations with the UTIP index than for the rest of the sample. This supports the stance that the imputation of missings in the underlying sectoral wage and employment data does not lead to systematically different numbers in the resulting inequality index. Second, in many cases with low correlations, the deviations between the UTIP and the new index are equally high across all versions of the index - that is the long, short, and full ones - making it fairly certain that the data used for the UTIP index again stem, at least partly, from other versions of the UNDIO industrial statistics.[24] If anything, correlations are lower with the short version of the index, which is an indication that the UTIP in some cases appears to also use only a subset of sectors for the calculation of their index. This is, presumably, a reflection of efforts to keep the measure time-consistent.[25] The

[22]The correlation is equal to one in Burundi, Benin, Burkina Faso, Belize, Congo, Cuba, the Dominican Republic, Gabon, Iraq, Kazakhstan, Kuwait, Nigeria, Puerto Rico, El Salvador, and Tanzania. Five more countries display negative correlations in differences: Australia, Belgium, Moldova, the Netherlands, and Puerto Rico.

[23]As measured by the total number of imputed values over all years and sectors for a given country. Appendix table 3.A.3 provides an overview of the correlation between the two measures and the extent of imputation.

[24]E.g., in Puerto Rico, Estonia, Bulgaria, Jamaica, and Uganda, among others.

[25]This is more prevalent for the levels than for the differences, which is in line with the previous finding that even when inequality levels are different, a slimmer version of the index is still able to trace changes over time quite well. By construction, countries where this was the case have been included in the "dynamic" version of the index and hence the higher deviation of the "short" version for the levels as compared to the differences is implied. The lower similarity with the UTIP also shows up in the average correlation across all countries, which drops to 0.6. To name a few country cases, lower correlations for the short

lower correlation with the short- as compared to the long version of the index occurs in several instances where the short version was kept due to the inaccurate representation of inequality levels and/or dynamics of the long version, as explained in section 3.2.3. This finding potentially casts doubt on the capability of the UTIP index to be a fully accurate indicator of manufacturing wage inequality. Lastly, in several instances, the deviations from the UTIP are substantially lower (but never zero) with the "full" (time-inconsistent) version of the Theil index.[26] Because the "full" index is based on a differential sectoral coverage over the years, this suggests that some of the UTIP numbers might rely on unbalanced underlying sectoral data, which would put into question the time consistency of their index. It could, however, again also stem from harmonization efforts with the previous version of the index.

In order to get an idea of the drivers of the divergence between the UTIP measure and the new index, a simple panel regression[27] is employed with the percentage difference between the UTIP- and the new index as the dependent variable. The number of dropped sectors and the number of imputed data points in the underlying sectors in each year are the main explanatory variables, and year dummies are added to the model to check whether the difference between the indices is growing over time. Table 3.2 contains the results.

Clearly, the number of imputations is related to the divergence of the two measures, with one additional imputed data point implying a 3 percentage point higher deviation. Interestingly, the number of dropped sectors has a negative coefficient, meaning that for every dropped sector, the two indices are on average 5 percentage points more similar. However, the use of robust standard errors, as warranted by a maximum likelihood ratio test, renders the coefficient insignificant. The last two columns do not contain the year fixed effects, which clearly reduces the size of the coefficient on the dropped sectors. Looking at the values of the year dummies (displayed in the full version of the table in the appendix, table 3.A.4, it becomes clear why this is the case: from 1990 onwards, the year dummies become positive and keep increasing over the 1990s and 2000s. This can be explained by the fact that, as mentioned in section 3.2.3, data for five sectors (19, 30, 32, 35, and 37) are only available from 1990 onwards and often only start in the mid-1990s. As explained earlier, they are therefore frequently dropped for the long version of the Theil index, and the same appears to be the case for the UTIP index. Additionally, harmonization efforts of the new and old UTIP index mainly focus on the

version are found in the Netherlands, Great Britain, Bolivia, and Romania, among many others. It should not go unmentioned that the opposite case also occurs in the data a few times, e.g., in Botswana, where the correlations jump from 0.23 from the long/dynamic to 0.95 for the short version in differences, or Madagascar, where they rise from 0.66 to 0.96 for the levels. It should be noted, however, that these cases also display almost equally high (and sometimes even higher, as, e.g., in Ireland,) correlations with the "full" (time-inconsistent) version of the index and the short version may merely be a reflection of sectoral coverage in the full version, especially if there are few years with missing sectors.

[26] Most notably, this is the case for Madagascar, New Zealand, Moldova, Great Britain, and Austria for inequality levels. The problem is less prevalent for differences, where the correlation is often higher with the short version than the "full" one.

[27] The initial idea was to estimate the model in fixed effects to account for the fact that the UTIP relies on data sources other than the UNIDO industrial statistics in some countries. However, a Hausman test indicates that the estimates do not differ from the more efficient random effects model (chi2(47) =49.61, p = 0.3697), which is therefore retained.

Table 3.2: Explaining differences to the UTIP index: imputation vs. sectoral coverage

	(1)	(2) r	(3)	(4) r
Imputations	3.030***	3.030**	3.111***	3.111***
	(0.466)	(1.225)	(0.460)	(1.154)
Dropped sectors	-5.282***	-5.282	-2.396***	-2.396
	(0.815)	(3.444)	(0.607)	(1.817)
Constant	0.0131	0.0131	0.105	0.105
	(10.30)	(5.216)	(4.026)	(4.081)
Observations	3.627	3.627	3.627	3.627
Year dummies	YES	YES	NO	NO
# of countries	135	135	135	135
R^2	0,036	0,036	0,016	0,016

Notes. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the percentage deviation between the dynamic version of the newly constructed Theil index and the UTIP index. All estimations are employing random effects. The "r" in the top column indicates that standard errors are robust. R^2 refers to the within R^2.

time periods in which the indices overlap, and data from older classifications schemes can only be employed before the transition from the old to the new classification took place. The year dummies pick up this effect which is similar across all countries and allow the coefficient on the sectoral coverage to capture the remaining variation in sectoral coverage. Also note that the constant is not significantly different from zero, which means there does not seem to be an inherent difference between the two indices.

Overall, while I can replicate a major part of the UTIP inequality statistics with the newly computed index, there are large differences in quite a few cases. This is in line with the fact that the explanatory power of the model analyzing the differences between the new index and the UTIP one is very low. These results suggest that other factors not contained in the model - in line with the descriptive evidence discussed above, one of them most likely being the underlying data source - are more relevant for causing the difference between the two indices. For the remainder of this paper, this implies that all conclusions drawn only apply to the inequality numbers based on the sectoral information from the UNIDO industrial statistics using the ISIC Rev. 3, and not necessarily to those stemming from other, possibly more detailed data sources or sources using other industrial classification schemes. However, given that future values will be in the new classification scheme as well, the relevance of my results will be growing as the time coverage of the index is extended to more recent years.

3.4 On the role of within-sectoral inequality

Although the UNIDO industrial statistics do not contain individual-level data, one can still compute part of the within-sectoral inequality by exploiting the more refined sectoral classifications up to the 4-digit level, as provided by the Industrial Statistics Database (INDSTAT4). The share in total wage inequality of the within-component at the 3- and 4-digit level can give at least a rough idea of the lower bound of overall manufacturing wage inequality.[28] Unfortunately, the time coverage is much lower than for the 2-digit level data and spans only for the years from 1990 onwards. It should be noted that the raw data at the 3- and 4-digit level suffer from the same problems of unbalancedness as the 2-digit ones, but have not been modified in any way to address the resulting problems of comparability.[29] Inequality numbers - both between- and within-sectors - are therefore not generally comparable over time and shall merely provide an indication of the potential magnitude of within-sectoral inequality.

Within-sectoral inequality is created at three levels, 4d representing the most detailed (4 digit) level. The formula, introduced in section 3.2.1, is, in its expanded version, easily separable into different components:

$$T' = \sum_{s=1}^{S} y_s \cdot ln(\frac{y_s}{n_s}) + \sum_{s=1}^{S} y_s \sum_{s3d=1}^{S3d} y_s 3d \cdot ln(\frac{y_{s3d}}{n_{s3d}}) + \sum_{s=1}^{S} y_s \sum_{s3d=1}^{S3d} y_s 3d \sum_{s4d=1}^{S4d} y_s 4d \cdot ln(\frac{y_{s4d}}{n_{s4d}})$$

The different parts are calculated separately in order to enable statements about the contribution of 3-digit "between sector"-inequality as the within-sectoral component at the 2-digit level, without adding the 4-digit level contribution as well. The following terms

[28]Note that the sum of the between-component and the within-components at the 3- and 4-digit levels is in the following referred to as "total" or "overall" inequality for the sake of simplicity, although it is technically not total or overall inequality given that within-sectoral inequality at the 4-digit level remains unaccounted for.

[29]Another problem of the multi-level data for the calculation of a decomposable Theil index is that subgroups must be exhaustive and mutually exclusive. I.e., all lower-level (4- and 3-digit) numbers must add up to the total value provided for next level. Since this is a necessary requirement, the raw data had to be adjusted in a way such that the numbers add up at the different levels. If the higher-level value was higher or lower than the sum of the lower-level values, the difference has been added to or subtracted from the higher-level figure. While a desirable alternative would have been to create an extra category at the lower level containing the missing amounts in the case of too-low sublevel sum, this would have meant that in some cases, positive numbers for one variable (wages or employees) are matched up with zeros for the other one, and including this "residual" sector in the calculation of the Theil index is impossible due to the logarithmic transformation of the ratios.

are retained separately:

Between sectoral inequality at the 2 digit level: $$BE2 = \sum_{s=1}^{S} y_s \cdot ln(\frac{y_s}{n_s})$$

Between sectoral inequality at the 3 digit level: $$BE3 = \sum_{s3d=1}^{S3d} y_s3d \cdot ln(\frac{y_{s3d}}{n_{s3d}})$$

Between sectoral inequality at the 4 digit level: $$BE4 = \sum_{s4d=1}^{S4d} y_s4d \cdot ln(\frac{y_{s4d}}{n_{s4d}})$$

Within sectoral inequality at the 3 digit level: $$WI3 = \sum_{s3d=1}^{S3d} y_s3d \cdot BE4$$

Within sectoral inequality at the 2 digit level

(without the 4 digit level contribution): $$WI2 = \sum_{s=1}^{S} y_s \cdot BE3$$

"Total" within sectoral inequality amounts to: $$WI2 = \sum_{s=1}^{S} y_s \cdot WI3$$

The final index is then computed as BE2 + WI. The average contribution across all sectors, countries, and years of within-sectoral inequality at the 3- and 4-digit levels (WI) is 33.7%, which would indicate that between-sectoral inequality (BE2) still explains around two thirds of overall inequality in manufacturing. In terms of contributions to the within-component of the 3- vs. the 4-digit level, interestingly, the one-third/two-third ratio found previously for the between- versus within 2-digit level is reversed. On average, little over one third of within-sectoral inequality stems from inequality at the more aggregate 3-digit level (BE3) while two thirds can be attributed to inequality between 4-digit level sectors (BE4). Of course, true total within-sectoral inequality will be larger given that inequality within the 4-digit level sectors remains unaccounted for here.

Especially the result for the overall contribution of the within-component should, however, be taken with caution given the unbalancedness of the raw data.[30] The actual 3- and 4-digit within-sectoral inequality is certain to be higher in years with larger gaps and more missing data at the lower levels, and a first, crude correlation analysis indeed confirms a small positive correlation of 0.2 between the number of subsectors per 2-digit category and the share of the within-component. Moreover, the variation in the importance of the within-component across countries and years is very large and there are cases where within-sectoral inequality explains as much as 87% (Moldova in 2002) of overall inequality. Country averages also show a lot of variation and range from 66.6% in Lebanon to 5.4% in Kuwait. There are no clear trends in the development over time, either - in some countries, the within component seems to be growing, in other it is decreasing, and in several cases it is relatively constant over the years. Again, it is important to keep in mind that at least part of the variation in the within component stems from the unbalancedness of

[30]This is less of a problem for subsectors at the 3- and 4-digit level, given that a missing 2-digit sector implies that all of its subsectors are missing as well, whereas a missing 3-digit sector "only" leads to missings at the 4-digit level, which is the smallest available bracket already.

sectoral coverage over the years. Appendix figure 3.A.2 displays, for every country, the development of both the contribution of the within component - that is, the percent of total inequality which stems from the within-component - and the total number of sub-sectors (both 3- and 4-digit) with non-missing values in all 2-digit categories per year. As can be seen in the graphs, there are several countries with consistently high (or low) data coverage, which potentially mask the importance of the balancedness issue for teasing out the true within-component. Indeed, the standard deviation of the variation in sub-sector coverage is a mediating variable in the association of the contribution of the within component and sectoral coverage.[31] Once the countries with a low variation in sectoral coverage are discarded, the correlation between the share of within-sectoral inequality and sectoral coverage rises to 0.25 (countries above the mean variation). Only using countries with one standard deviation above the mean variation, it is 0.52, and for countries with variation higher than two standard deviations above the mean, it is 0.71. This indicates that very large changes in sectoral coverage are accompanied by increases in the importance of the within-component as well. Nevertheless, balancedness does not seem to be the only driver of the importance of the within component across the entire sample, and in particular it is not very relevant for those countries with good data coverage throughout.[32]

The second major factor for the extent of within-sectoral inequality is the sectoral composition of the manufacturing industry of a country. Some sectors by construction have more subsectors than others. In the extreme case of only one subcategory per 2- and 3-digit category, there is no within-group inequality by construction. This is the case for sectors 16 (Tobacco products) and 30 (Office, accounting and computing machinery), and consequently, countries whose manufacturing industry is concentrated in those sectors are likely to have a lower share of within sectoral inequality. Averages across sectors indeed reveal large differences in the importance of the within-component, and the ranking of 2-digit sectors in terms of the size of their within-component (taking both 3- and 4-digit sectoral inequality into account) is clearly correlated with the number of subsectors into which each category is divided.[33] Appendix table 3.A.5 provides more detailed information on the association between the number of subsectors and the size of the within component for every sector.[34]

In order to work out the importance of the sectoral composition, a simple country and year fixed effects regression is conducted, where the share of within-sectoral inequality is

[31] The correlation between the standard deviation in the total number of subsectors (across years within a country) and the correlation of the same with the share of the within component is 0.44.

[32] Another explanation for the low average correlation is the very crude measure of data coverage provided by the total number of subsectors. It could still very well be - in fact, it is likely to be the case fairly often - that some sectors are included in some years while others are not. This variability is very likely to substantially affect the within-component. In order words, it does not only matter how many sectors are included, but also which ones are included (and which ones are not).

[33] The correlation is 0.75 and the number of subsectors refers to the mean number of 4-digit sectors per 2-digit category. The correlation with the total number of subsectors (3- and 4-digit sectors) is very similar (0.77). Only cases which have a non-zero within-component have been considered in the calculations.

[34] Clearly, those sectors ranking high on within-sectoral inequality (that is, the logged ratio of the wage- over the employment share) also tend to have a higher number of subsectors. This is still true for the weighted component shown in panel 3 of table 3.A.5, although the association is slightly weaker due to the weighting with the sector's wage share shown in panel 2.

regressed on the number of subsectors covered by the data and a set of year dummies.[35] The results are displayed in column 1 of table 3.3. For an assessment of the importance of sectoral coverage versus sectoral composition, the wage shares of the different 2-digit sectors are then added to the regression in column 2.[36]

Clearly, the sectoral composition takes away from the sectoral coverage effect, which decreases by more than 40%. While the fixed effects estimator does remove all time-invariant country-specific factors which potentially affect the size of the within-component, it still estimates a common slope parameter for the sectoral coverage variable for both high- and low-variability countries. Random effects estimation confirms that the coefficient on the sectoral coverage variable ("subsectors") is hardly affected by the removal of the country fixed effects (as shown in appendix table 3.B.6). As argued above, sectoral coverage is likely to be a relevant factor skewing the size of the within-component only for those countries where sectoral coverage varies substantially over the years. The fixed effects regression is therefore repeated for two different high-variation subsamples: one with above-average variation in sectoral coverage, and one with one standard deviation above the average variation. Table 3.3 displays the results.

According to the point estimate of the number of subsectors, an additional sector is associated with a 0.026 percentage point higher wage share. Although it is highly significant and robust across specifications, this is a rather small number. Relating it to the standard deviation of the "subsectors" variable, an increase of one standard deviation (43) would imply a mere 1.12 percent higher within-sectoral wage share. For the high-variation subsamples, the point estimates rise to 0.03 and 0.05, implying higher within-sectoral wage shares of 1.4 and 2.1 percent, respectively, for a one-standard deviation increase in subsectoral coverage. If one considers the average maximum distance between the highest and the lowest sectoral coverage within a country, numbers for between-sectoral inequality would be between 3 and 5.75 percent higher on average. More noticeable than the increase in the coefficient is the substantial rise in the R-squared for the high-variation subsamples. It does indeed seem that sectoral coverage explains the lion's share of the variation in the within-component in those countries displaying major changes in sectoral coverage.

In order to get at least a rough idea of what the within-component would be if sectoral coverage had been larger in those countries displaying large variations over the years, the coefficient estimates obtained from the above regressions are used to obtain counterfactual

[35]The results presented use the number of 3-digit categories per 2-digit sectors because, as previously show, the 3-digit level accounts for two thirds of the within-component. Results are very similar when the number of 4-digit categories, or the number of total subcategories is used instead (results available upon request).

[36]Note that sectors 16 and 30 have been omitted from the regressions. If all sectors (including 16 and 30, which have no subsectors and can therefore never positively contribute to the share of the within-component) are included in the regression containing the sector shares, all sectoral coefficients have positive signs and interpretation of the results is not straightforward. This is because the shares of all other sectors are implicitly evaluated against the shares of sectors 16 and 30, which by definition (due to the lack of subsectors) never contribute to within-sectoral inequality. Hence, the larger the shares of the other sectors, the smaller will be by construction the share of sectors 16 and 30, which ceteris paribus implies a larger within-component. Moreover, because both sectors 16 and 30 have on average larger wage- than employment shares (the ratios being 176 and 137, see table 3.1) , whenever the wage share of those sectors rises, the between component of the Theil index will rise as well, implying by definition a smaller within-component.

Table 3.3: Within-component: sectoral coverage vs. sectoral composition, FE results

	(1)	(2)	(3) High variation sample1	(4) High variation sample2
Subsectors	0.0451***	0.0261***	0.0301***	0.0505***
	(0.016)	(0.009)	(0.009)	(0.012)
Share_15		0.799	1.805**	-0.205
		(0.970)	(0.798)	(2.772)
Share_17		-0.0140	0.728	-0.797
		(0.950)	(0.838)	(2.224)
Share_18		-0.616	-0.360	-5.747**
		(0.951)	(0.915)	(2.361)
Share_19		-0.654	1.798	8.376**
		(1.306)	(1.732)	(3.015)
Share_20		2.989**	4.801***	6.038*
		(1.416)	(1.102)	(2.934)
Share_21		0.370	1.637	1.527
		(1.401)	(1.186)	(3.452)
Share_22		-0.452	-0.466	-6.338**
		(0.967)	(1.015)	(2.013)
Share_23		-1.629	-0.826	-3.250
		(1.508)	(0.580)	(2.187)
Share_24		-0.542	-0.409	-2.067
		(0.976)	(0.879)	(2.292)
Share_25		0.604	0.848	-3.434
		(1.037)	(0.992)	(1.965)
Share_26		1.135	1.419	-0.253
		(1.040)	(1.069)	(2.192)
Share_27		-1.132	-0.494	-3.148
		(0.817)	(0.725)	(2.401)
Share_28		0.658	1.135	-0.683
		(0.866)	(0.836)	(2.440)
Share_29		-0.181	0.358	-0.784
		(0.900)	(0.833)	(3.053)
Share_31		-0.489	0.119	-2.919
		(1.231)	(0.852)	(3.596)
Share_32		-0.578	-0.00135	-4.494*
		(0.902)	(0.947)	(2.454)
Share_33		-0.108	-0.685	-6.274
		(1.516)	(2.246)	(4.080)
Share_34		-0.446	2.191*	-8.611
		(1.028)	(1.109)	(5.839)
Share_35		-0.567	0.124	-1.805
		(0.877)	(0.838)	(2.448)
Share_36		-2.478*	-1.338	-4.480
		(1.297)	(1.311)	(2.993)
Share_37		2.189	1.967	2.196
		(2.656)	(2.070)	(6.554)
Constant	15.14**	26.44	-46.73	201.4
	(6.63)	(85.10)	(73.66)	(223.0)
Year FE	YES	YES	YES	YES
Observations	429	429	221	74
R-squared	0,465	0,465	0,654	0,898
# of countries	53	53	27	11

Notes. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the share of within-sectoral inequality in %. Numbers 15 to 37 refer to the 2-digit sector's wage share in total manufacturing wages in %. High variation samples 1 and to refer to subsamples of countries with above-average variation in sectoral coverage (1), and countries with one standard deviation above the average variation (2).

values when sectoral coverage is raised to a level found in other years in the same country.[37] The adjustment is done for all three samples.[38] The first adjustment on the full sample yields a within-component of 34.5, which is less than 1 percentage point higher than the unadjusted value of 33.7. The adjustment for the first subsample, displaying an above-average standard deviation of sectoral coverage, results in a very similar value of 34.4 for the full sample. However, adjusting the high-variation subsample leads to a substantial increase in the average within-component to around 40.5%. Given that the rest of the sample remains unadjusted, the true extent of within-sectoral inequality at the 3- and 4-digit level is likely to be above 40%. Furthermore, sectoral composition has not yet been accounted for, either. As shown in table 3.3, including the sector shares takes away from the positive effect of sectoral coverage and hence sectoral composition is likely to increase the within-component even further. This is especially true for those countries with a lower variation in sectoral coverage where the sectoral coverage adjustment would make very little difference.

The full extent of within-sectoral inequality in manufacturing is of course not covered by including sectoral data at lower levels. There is certainly a substantial amount of inequality within 4-digit sectors, which on average still have almost 22,000 employees. In a country like Great Britain, where the within-component accounts for as much as 85% of overall inequality in 2009, average employee numbers in that year at the 4-digit levels are 44,600, leaving room for a substantial amount of unequal pay among these workers. If inequality within the 4-digit sectors was added to the within-component, it is fairly certain that between-sectoral inequality would explain very little of the overall inequality in manufacturing.

Nevertheless, it is possible that changes in inequality over time in between-sectoral inequality can reflect the overall trends in inequality. Conceição and Galbraith (2000: 67) argue that this is likely to be the case, given that "while within-group inequalities are likely to be large relative to differences between group averages, the internal rigidity of industrial structure tends to assure that changes in within group inequalities in an industrial classification will be small relative to changes between groups." However, building on the argument that "industries [...] mean something, and if they mean anything at all, the effect must be to impose a measure of homogeneity on entities classified together, and a measure of distinctiveness to entities classified as being in different groups," it is much

[37] It is not obvious what this level should be and setting it is somewhat arbitrary. In order to not over-estimate the potential within-sectoral inequality numbers due to outliers at the top, the sectoral coverage numbers are split into quintiles and the lowest value of the highest quintile is used as the counterfactual coverage value for all years with lower sectoral coverage. When numbers of sectoral coverage are identical at the upper end of the distribution, or when there are too few data points for a country, the highest value of the 4th quintile is used instead. In those cases where both values are available, the difference between the two is very small (8 on average), with a maximum of 33 for Ireland, where the values are rather high at 415 and 448.

[38] Because sectoral composition is by construction skewed in those years with missing data because sectoral shares for non-missing sectors are larger than they would be if the missing sectors were present, simply using fitted values from the previous regression model would distort the results substantially for precisely those years where sectoral coverage is lower. Therefore, only the coefficient estimate for sectoral coverage is used and is multiplied with the yearly difference between the counterfactual high data coverage and the actual number of 3-digit level subsectors. This value is then added onto the observed within-component. Where the counterfactual high data coverage is lower than the actual number of 3-digit level subsectors, the original value is retained and consequently, the within-component is not modified.

easier for a worker to switch between subsectors within a broad industrial category than to switch between industries. In fact, it is possible that there is no change in inequality at the broader 2-digit sector level, with employee and wage bills remaining unchanged, but there can be a substantial amount of re-shuffling within industries which remains unaccounted for entirely. It is true that for a large number of sectors (wherein 22, being the number of manufacturing sectors covered by the data, can be considered large), the overall effect of large within-changes is mitigated due to the presence of 21 other sectors. However, these 22 sectors are far from equally sized, and it is precisely the large sectors which are divided into more subcategories and display larger amounts of within-sectoral inequality to begin with. Most notably, sector 15 (food and beverages) makes up 20% of the wage share on average. Looking at developing countries, where this sector is of higher economic significance than in the developed world, it accounts for over one fourth of overall manufacturing wages. It also has the highest within-sectoral inequality, and, consequently, more than forty percent of within-sectoral inequality can be attributed to sector 15 on average in developing countries. The average contribution of 14% of the sector to between-sectoral inequality at the 2-digit level is also large. It is, however, to a large extent attributable to the sector's large wage share - in fact, sector 15's discrepancy between wage- and employment shares is among the lowest of all sectors, at least for developing countries. Hypothetically, if one assumed an increase in wages in sector 15, this would decrease between-sectoral inequality because the sector's contribution to the Theil index is negative, i.e., it has a lower wage- than employment share. Nevertheless, the within-sector component would be assigned a higher weight due to the sector's increased wage share, even assuming that the increase in wages is distributed within the sector in such a way that does not lead to higher within-sectoral inequality itself. As a result, between-sectoral inequality would decrease but within-sectoral inequality would increase. Unfortunately, the unbalancedness of the 3- and 4-digit level data, which is even more severe for developing countries, makes it difficult to empirically test whether this scenario is occurring in practice.

What is feasible, however, is a check of whether the data given in any single year would theoretically allow for this case to happen. That is, the change in the wage share which would lead to a zero within-component for a given sector is multiplied with the within-component (which, for a conservative scenario, is assumed to remain unchanged). This increase in the within-component is then compared to the maximum possible inequality decrease in the opposite direction for the between-sectoral component. Mathematically, this is equivalent to comparing the following two elements in a simplified two-sector scenario:

$$\frac{\Delta Y_j}{Y}T_j = \frac{Y_j}{Y}ln(\frac{\frac{Y_j}{Y}}{\frac{n_j}{n}}) + \frac{1-Y_j}{Y}ln(\frac{\frac{1-Y_j}{Y}}{\frac{1-n_j}{n}})$$

The first element is made up of the difference in the new and the old wage share of the sector of interest j, which is the weight for Tj, representing inequality within the sector. The second element is the between-component. One can think of this as a hypothetical 2-sector scenario in which all other sectors (which, for simplification purposes, are assumed to not display any within-sectoral wage inequality) apart from the sector of interest are aggregated into one large sector. Although the interest here is in comparing changes in the two elements, because the minimum value for between sectoral inequality is zero, the maximum possible decrease is equal to the entire between-sectoral component at the 2-digit level - arguably, a rather unrealistic scenario, but nevertheless one which serves in proving the point that within-sectoral inequality trends can outweigh between-sectoral movements in inequality. What is not so unlikely is an increase in wages to a level that leads to a contribution of a large sector close to zero, given that the employment and wage shares are already relatively equal for sector 15 in many cases.

It turns out that there is only one single case in the data where in the above scenario it is theoretically possible that the first component outweighs the second, namely sector 15 in Rwanda in 2009. It is worth noting, however, that in a number of cases, the two effects - the inequality-decreasing effect of the between-component and the inequality-increasing effect of the within-component - almost cancel themselves out. The true decrease in overall wage inequality is therefore substantially lower than what it seems if only the change in the between-component is considered and the opposing effects of the within-component are ignored.

There are more reasons to believe that the between-component is a poor indicator of overall movements of between-sectoral inequality. First, it seems implausible that overall wage inequality drops to zero as a result of an increase in wages in one sector. The true decrease in the between-component in the above scenario is therefore likely to be much smaller, leaving more room for the within-component to counteract this effect. Second, the within-component is still vastly underestimated in many countries and years due to the unbalancedness in the raw data. This may be one of the reasons for why the above counterfactual exercise only yields a single case in which the within-component could outweigh the between-sectoral effect if the latter drops to zero. Third, adding to this underestimation, the within component is in all cases missing a further element due to the lack of individual-level data. Fourth, the assumption of a zero change in within-sectoral wage inequality was made to demonstrate the most conservative (and mathematically most simple) case of changes in the two components, where the change in the within-component was constructed to be minimal and the change in the between-component to be maximal. If the assumption of a zero change of the within-component is dropped as well and replaced with an increase in within-sectoral inequality - which is, after all, the scenario we are truly interested in - it is very likely that more cases can be identified in the dataset which have the potential to display divergent trends in between- and within-sectoral wage inequality. Going through different scenarios of changes in the within- and between-components is a tedious exercise, which can be circumvented by directly performing comparisons of the two components on the raw data. Despite the previously discussed limitations of comparing changes over time due to the unbalancedness of the 3- and 4-digit level data, conclusions can still be drawn from comparisons of the direction of changes of the within- and the between component given the following considerations. While it is clear that the size (and hence the share in overall inequality) of the within-component is affected by the availability and composition of subsectoral data, this does not affect the change observed in the between component. Assuming that the pattern of missings is random across subsectors and it is not the case that sectors with higher within sectoral inequality are missing more (or less) often than those with lower sectoral inequality, looking at changes between years with similar subsectoral coverage can provide some indication of whether changes in the between- and the within-component go into different directions. Indeed, out of the 968 observations with changes of less than 1% in subsectoral coverage from one year to the next, around 13% (stemming from both developing and developed countries from all regions) show opposing trends of changes in overall and changes in between-sectoral coverage. The change in the within-component goes in the opposite direction than that of the between-component, and is large enough to outweigh its effect. This picture changes very little if only those cases are considered where sectoral coverage is entirely unchanged: again, around 14% of the 294 cases show different trends in overall and between-sectoral wage inequality.[39]

In sum, there is strong indication that the Theil index relying on the between-sectoral component of manufacturing wages computed here and by the UTIP may provide a wrong image of overall changes in manufacturing wage inequality in around 13% of cases. Given that the "true" extent of within-sectoral inequality (taking into account individual-level data) is likely to be substantially larger, this number has to be taken as a lower bound to the true discrepancy between between-sectoral and overall changes in manufacturing wage inequality.

Of course, many more things can be done to assess the plausibility and extent of error of only looking at between-sectoral changes in inequality. Besides the counterfactual exercises on the UNIDO data discussed above, one could look into country cases with better data for manufacturing wages. This would allow, at least in some cases, the calculation of inequality up to the individual level and provide some indication of the remaining extent of inequality not captured by the sector-level data, no matter how detailed. However, doubts also arise about "external validity" of the index in the remainder of this paper, which, if taken seriously, limits its relevance to a narrowly defined set of applications focusing specifically at manufacturing. I therefore leave it up to those who have such a confined focus and need to take into account changes within sectors to assess this last component of within-sectoral inequality which remains unaccounted for here.

3.5 The relationship to overall income inequality

3.5.1 Comparison with overall inequality statistics

In an effort to validate the capacity of the Theil index to serve as a proxy for, and basis of, developments in overall monetary inequality, Galbraith and Kum (2005) (henceforth GK2005) relate it to the Gini coefficients compiled by Deininger and Squire (1996) (DS1996). They find an elasticity of between 6 and 8.5 %, which they explain with "the much greater volatility of the Theil measure [due to the varying number of manufacturing industries per year and country], and also the greater volatility of manufacturing pay

[39]The result also holds for different threshold of change in subsectoral coverage of between 5% and 50%. In fact, the share of 13% remains remarkably stable across all chosen thresholds.

compared with household income [because it includes income from other sources such as non-labor wage, land and capital]" (GK2005: 128). Adding the share of employment in manufacturing as a control variable, and dummies for the different income categories underlying the DS1996 Gini coefficients, they use the predicted relationship between their Theil index and the Gini coefficients to scale up the Theil index and obtain a broad measure of income inequality, the "Estimated Household Income Inequality" (EHII) dataset. In a more recent update of their estimates, they confirm the relationship with the DS1996 data (Galbraith et al. 2014, 2015).

This paper aims to expand upon this approach and to also explore the circumstances under which the two measures are more (dis-)similar by regressing the difference between the Theil index and overall income inequality on a number of control variables. Instead of the Deininger and Squire (1996) dataset, which only contains data until 1996, the Gini coefficients from the World Income Inquality Database (WIID) provided by UNU-WIDER are used which comprise and extend the DS1996 data. All the controls proposed by GK05 are included and several other potentially important determinants of the association between manufacturing wage inequality and overall inequality are added to the model. To eliminate volatility stemming from differential sectoral coverage of the Theil index, the normalized version of the Theil index is used so that numbers are comparable between countries with differential sectoral coverage. To tackle another potential source of volatility in the income inequality measure, a few more control variables are added, and other, arguably more consistent, measures of income inequality are used in addition to the WIID.[40]

Most notably, the Luxembourg Income Study (LIS) Gini coefficients are derived from harmonized primary micro data, which is currently considered the "gold standard" in terms of consistency and accuracy of the resulting inequality measures. Unfortunately, time and country coverage of the LIS data is still limited and the resulting sample size is correspondingly small. Nevertheless, the comparison with the LIS can be used to validate the results from the WIID. The second alternative measure of income inequality used are the SWIID Gini coefficients provided by Solt (2015). He addresses the inconsistencies in the WIID, arising from the previously discussed heterogeneity of data sources by providing multiple solutions to mitigate the same and combining them into a single workable dataset.[41] The result is a balanced multiply imputed dataset of broad country and time coverage. Although the underlying method has been criticized (Jenkins 2015), the SWIID certainly provides a more sophisticated, prudent, and explicit way of making the WIID data comparable, especially when compared to the much more crude alternative of merely introducing dummy variables for the numerous categories of income and other underlying

[40] It should not go unmentioned that the creators of the EHII have put their resulting estimates through a number of validity checks and comparisons with other data on income inequality, including the LIS (Galbraith et al. 2014, and 2015). They have not, however, repeated the initial exercise of relating the different data sources directly to the UTIP-UNIDO index of wage inequality.

[41] Importantly, it should be noted that the SWIID uses other data sources to cross-check its values, among them the UTIP-UNIDO Theil index. One might therefore suspect a built-in association between the index calculated in this paper and the SWIID Gini coefficients which is closer than for the other data sources. As shown in table 3.4, this is clearly not the case and suggests that the use of the SWIID is unproblematic in this context.

concepts on which the WIID relies. That the latter approach imposes constant differences between concepts across countries and over time is just one of its problems[42] and has been shown to be invalid (Atkinson and Brandolini 2009, Galbraith and Kum 2003. The SWIID data have another advantage: they provide Gini coefficients for both market and net inequality, and thereby render the inclusion of a variable for government transfers, as suggested to GK2005 but excluded by the authors due to insufficient data coverage of the variable, unnecessary. Manufacturing wage inequality is expected to show a closer relationship with market- than with net inequality since the latter includes transfers which are designed for the very purpose of mitigating inequalities arising from (inter alia) wages. Two further data sources are added: the EU SILC data and the Gini coefficients from the World Development Indicators (WDI). They are tested here to allow comparisons to Galbraith et al. (2015), who use them to validate their EHII data.

The three control variables used in GK2005 are the ratio of manufacturing employment to population, the share of urban population, and the population growth rate. Apart from measuring the importance of manufacturing for overall incomes, the share of manufacturing is also supposed to capture the part of the labor market which tends to be more unionized, and is therefore expected to be associated with lower inequality. Urbanization is expected to be associated with more inequality because "wealthy people live in cities," and population growth serves as a proxy for the age structure of a country and the composition of households, and is expected to be associated with higher inequality at the household level if poorer households tend to be larger.

Instead of the ratio of manufacturing employment, the share of value-added in GDP from manufacturing (Mfg_va) is used here, which features a good coverage of the countries in the sample.[43] The variable, along with the share of urban population (urban) and population growth (gpop) is taken from the WDI (2016). In addition to these variables, three more controls are added. The first one is the price level of investment (pl_i), taken from the Penn World Tables (PWT, V8.1, Feenstra et al. 2015). The variable is a proxy for the rate of returns of capital, and since capital is a component of overall income, higher returns to capital might increase the divergence of the two measures.

GK2005 argue that one of the reasons why changes in manufacturing wage inequality will likely not counteract developments in overall wage inequality is that low-skilled workers, forming the lower end of the distribution in manufacturing wage inequality, are substitutes for low-skilled workers in other sectors such as agriculture and services. It is therefore unlikely that wages at the lower end of manufacturing pay decrease or increase without an equivalent shift in the wage levels of other sectors of the economy. That being said, the same logic does not apply to the upper end of the wage spectrum, where workers are skilled in a specific profession and are much less likely to easily switch between manufacturing and other sectors. To also account for changes at the upper end of the wage

[42] Another problem is how to deal with multiple observations per country and year of the same quality. Here, the researcher faces a trade-off between various dimension, e.g., sacrificing demographic for geographic coverage. Approaches which directly adjust the WIID Ginis by adding or subtracting the average differences between the underlying categories, as e.g. in Gruen and Klasen (2012) and Easterly (2007) do not circumvent the problem, either, since differences remain also for the adjusted Ginis.

[43] The variable leads to very similar coefficient estimates as found in Galbraith and Kum (2005) when their index is used on the WIID data. Results can be found in appendix table 3.A.12.

distribution, a measure of total factor productivity (tfp) is included which is constructed to reflect cross-country differences in aggregate technology (Feenstra et al. 2015). As the technological frontier of a country shifts outwards, this is likely to encompass all sectors of the economy and affect skill premia everywhere. Technological change can therefore be assumed to sway manufacturing wage inequality and overall income inequality in the same (upward) direction.

Naturally, both of these mechanisms are weakened by the extent of openness of an economy. While "cheaper" foreign workers may - indirectly through trade - be substitutes for some low-skilled manufactures, the same cannot be said for non-tradable services or some segments of agriculture. Similarly, countries can gain access to technology through trade, which may again affect the tradable sectors more than the non-tradable ones. In addition to these effects, trade openness (defined as import and export value over GDP and taken from the WDI) is also a proxy of the extent to which a country is vulnerable to external shocks which are likely to affect overall (income) inequality much more and cause divergences from the wage inequality measure. To generally account for shocks which potentially affect all countries, year dummies are added to the model as well. Lastly, since many of the above control variables are correlated with GDP per capita, it is included to make sure that its effect gets picked up separately.

In addition to these "external" variables, the number of imputations is added to the regression to account for the fact that in the case of linear interpolation, the idea was to be as conservative as possible in mapping observed changes in employment and wages in the underlying sectors over the missing years. Consequently, actual changes which may show up in the overall income inequality statistics are less likely to be captured in those cases where imputations were necessary. The impact of the second "internal" variable, sectoral coverage of the wage inequality measure (# of ISIC), is time-invariant and cannot be estimated with the fixed effects approach. Therefore, random effects estimations are employed additionally to get an idea of the role of this variable as well as of the impact of controlling for all other country-specific time-invariant factors.

Before moving to the analysis of the deviations between the two measures, the specification by GK2005 is replicated using my newly constructed index and the WIID instead of the DS1996 data. One the one hand, this serves as a check as to whether the new Theil index yields results similar to theirs. On the other hand, it tests whether their results also hold with the extensions discussed above which will be used in the analysis of the deviations of the measures. Having reduced the two sources of volatility identified by GK2005, the relationship between the inequality data and the Theil index should be stronger in general, and in particular with the newly added, more consistent income inequality measures. Finally, it gives a first idea of how the new control variables relate to overall (income) inequality. Table 3.4 contains the fixed effects results for the expanded specification of GK2005 and features the estimates from their paper in the first column for better comparison.[44] As in their model, the Theil index enters in logs to simplify

[44]Note that these are not based on estimations done in this paper, but they are literally the numbers published in table 5 of their paper.

interpretation and to account for its log-normality (as shown in appendix figure 3.A.3).[45] The random effects results are displayed in appendix table 3.A.8 and do not show major changes in the results. Note, however, that the Theil index should not be employed as a representation of *overall* manufacturing wage inequality in random effects models in general, given that within-sectoral is not accounted for and the measure hence massively understates overall manufacturing wage inequality levels. Fixed effects models on the other hand only consider mean-deviations over time and hence might still be able to trace changes in overall manufacturing wage inequality - although the accuracy thereof is questionable as well (see section 3.2).[46]

First, comparing columns 1 and 2 of table 3.4, the coefficients on the Theil index are substantially lower than those found by GK2005 for the UTIP measure. This is not due to the use of the DS1996 instead of the WIID data in column 1. Very similar coefficient estimates are obtained when the UTIP index is regressed on the WIID instead of the DS1996 data (results for the reduced model can be found in appendix table 3.A.12.) Despite the larger sample size, significance disappears in the fixed effects model, and drops to the 5 percent level in the random effects specification, as shown in column 2 of appendix table 3.A.8. The coefficient remains small and insignificant for all other inequality measures. Ironically, the only variable displaying similar effects as in the GK2005 estimations is the importance of the manufacturing sector, which is based on a measure different from theirs. The negative coefficients are more in line with the interpretation of the manufacturing sector as a proxy for the extent of unionization than as a mediating variable capturing the role of manufacturing wage inequality for overall income inequality, but are very small throughout.

While the results look rather similar for most measures, many coefficients change drastically when the LIS data are used. Most notably, the sign on the Theil index becomes negative (although the standard error is very large). The other variables population growth and urbanization also change substantially and are significant in some cases, despite the small sample size. When using only the LIS countries in the other specifications, it becomes clear that this is entirely due to the sample composition.[47] The fact that these differences arise between different samples despite the fact that country fixed effects are contained in the model also puts into question the universality of the relationship between the two measures for all countries.

Regarding the two versions of the SWIID, interestingly, the Theil index appears to

[45] In addition to the changes in the model described above, it also contains dummy variables for the underlying categories in the WIID data (for the full list, consult appendix table 3.A.7), while GK2005 only include dummies for the income concept and the income sharing unit (i.e., whether the data were measured at the household- or the person level) used in the DS1996 data. Note that the results do not change much when only the set of variables used in GK2005 is included (the fixed effects results can be found in appendix table 3.A.9). The largest change in coefficients is triggered by the inclusion of the GDP per capita variable, which affects the estimates of the control variables, but not that of the Theil index (results available upon request).

[46] Although the fixed effects model is clearly preferable due to the removal of time-invariant country specific factors, the random effects model is estimated to be able to compare also the random effects estimates by GK2005, and to get a benchmark estimate of the effects of the time-invariant factors on the income inequality measures to be able to better interpret the results from the next specification trying to explain the differences between the Theil index and the income inequality measures.

[47] Results available upon request.

be more closely related to net inequality than to market inequality, as indicated by the consistently larger coefficients on the former, although all are insignificant. One explanation for this finding is that it is not primarily labor market income of employees which is redistributed, and that redistributive taxes are on average similar across sectors so that the post-tax distribution of wages resembles the pre-tax one, but the post-tax distribution of overall incomes now has less resemblance to the pre-tax distribution of wages.

Table 3.4: Relationship between Theil and Income inequality: FE results, extended model

	(1) GK05 (DS1996)	(2) wiid	(3) swiid_n	(4) swiid_m	(5) lis	(6) silc	(7) wb
ln(Theil)	0.079***	0.0168	0.00776	0.00521	-0.00260	0.0272	0.00956
	(6.60)	(0.0123)	(0.0128)	(0.0136)	(0.0323)	(0.0166)	(0.0177)
GDPpc		-9.46e-07	1.42e-06	9.69e-07	-4.10e-06	-2.30e-06	1.23e-06
		(3.02e-06)	(2.02e-06)	(2.50e-06)	(3.58e-06)	(2.86e-06)	(5.35e-06)
Pop. growth	-0.578	-5.58e-05	0.00505	0.00727	0.0121	-0.0185	-0.00472
	(-0.81)	(0.0130)	(0.0101)	(0.0110)	(0.0347)	(0.0131)	(0.0169)
Share urban	0.001	-0.00950**	0.00206	-0.00148	0.0205*	-0.00365	0.00770**
	(-1.57)	(0.00408)	(0.00312)	(0.00374)	(0.0118)	(0.00512)	(0.00314)
Manuf.v.add.	-0.001***	-0.00683*	-0.00323*	-0.00430**	-0.00439	-0.00244	0.000585
	(4.50)	(0.00356)	(0.00188)	(0.00196)	(0.00773)	(0.00593)	(0.00334)
Trade open.		0.000746*	0.000148	2.53e-05	0.00101	0.000423	0.000440
		(0.000447)	(0.000310)	(0.000364)	(0.000802)	(0.000526)	(0.000494)
Price level inv.		0.0217	0.0463	0.0589	-0.116	0.107	0.0839
		(0.0559)	(0.0328)	(0.0366)	(0.0919)	(0.0920)	(0.0619)
# imputed		-0.00244**	0.000900	0.00118	-0.00306	0.00157	-0.00101
		(0.00120)	(0.00123)	(0.00156)	(0.00320)	(0.00152)	(0.00174)
Tfp		0.264**	0.0612	0.0780	0.329	0.321**	0.210*
		(0.110)	(0.0501)	(0.0719)	(0.248)	(0.117)	(0.113)
Constant	3.893***	3.923***	3.479***	3.887***	-2.838***	3.461***	2.991***
	(51.38)	(0.324)	(0.132)	(0.181)	(0.964)	(0.417)	(0.279)
Year dummies	NO	YES	YES	YES	YES	YES	YES
Observations	481	618	1.521	1.511	120	256	483
R-squared	unreported	0,81	0,12	0,166	0,522	0,25	0,188
# of countries	81	66	82	82	35	28	73

Notes. Robust standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1. The dependent variable is the logged Gini coefficient from the data source indicated in the top row. Swiid_n and swiid_m refer to net and market inequality, respectively. Silc denotes the EU SILC data and wb the WDI Gini coefficients. Instead of the share of manufacturing value added, the model by GK05 displayed in column 1 uses the ratio of manufacturing employment. It also uses the DS1996 data as the dependent variable, and the UTIP Theil index. Note that the coefficients displayed in column 1 are not based on estimations done in this paper, but they are literally the numbers published in table 5 of GK05. Their model contains a set of dummy variables for the underlying welfare concepts in the DS1996 data. Column 2 contains the full set of dummy variables for the underlying welfare concepts, income definitions, and other categories used in the WIID. The full results, including the year dummies, can be found in appendix table 3.A.7.

As for the newly added control variables, only total factor productivity has a stable positive effect across all models and is significant and rather sizeable in some cases. According to the WIID fixed effects point estimate, it increases income inequality by around 30%. The effect of GDP is negligible, but it is kept in the model because its inclusion affects the estimates of some of the other controls such as population growth, and urbanization.[48]

In terms of the "internal variables," the effect of the time-invariant variable sectoral

[48] Results available upon request.

coverage can be seen in the random effects results (table 3.A.8). The coefficient is small and only significant for the LIS data. This is reassuring given that the Theil index has been normalized with the underlying number of sectors for the random effects estimations already. It does not seem to be the case that countries with better sectoral coverage systematically differ from those with worse coverage, even when many time-invariant factors are not controlled for. This, again, supports the stance that the missings in the underlying sectoral data are random.

Overall, these results do not lend much support for the findings of GK2005, which suggest a stable association between the Theil index of manufacturing and overall inequality. While the findings for the WIID data are qualitatively still relatively similar to their estimates, other, and arguably more consistent, measures of income inequality yield rather different results. Not only does the association with the Theil index become insignificant, but the coefficients are also too small to postulate any economically meaningful link between the two variables. This is true even for the WIID specification, where according to the (insignificant) WIID fixed effects point estimate, doubling the Theil index would lead to an increase in the Gini coefficient of little under one percent. Given the high R-squared of 0.81 in the specification, the low association does not seem to be the result of an incomplete or widely misspecified model, either. One reason for the weak association of the between-sectoral Theil index and overall measures of income inequality could be the neglect of the within-component. It might be worthwhile to repeat the exercise with manufacturing wage inequality measures using more detailed sector- or individual-level data. That the UTIP index displays a stronger link with the income inequality measures could also be owing to the fact that it partly relies on earlier industrial classification schemes with higher levels of detail.

Given that there are good theoretical reasons to expect a robust relationship between manufacturing wage inequality and overall income inequality, an explicit analysis of the factors which might cause the two measures to differ stands to reason. All of the theoretically motivated variables discussed above are included in the model, along with the full set of year dummies and, for the WIID data, the underlying categories. The dependent variable is the logged percentage difference between the (normalized) Theil index and the respective Gini coefficient, as indicated in the top row of table 3.5.[49] The logarithmic transformation is used, on the one hand, to make interpretation easier, and on the other hand, because the differences are approximately log-normally distributed (see appendix figure 3.A.4).

With a few exceptions, the results do not match the theoretical predictions derived for the variables above, and most coefficients are insignificant. To begin with, a higher share of manufacturing value-added is associated with a higher discrepancy of the Theil index and income inequality. Apart from the smaller size on the WIID data, the coefficient is rather stable across the different data sources, and significant for the SWIID. In line with the interpretation that the variable is capturing the extent of unionization, one way of reading this result is that with a larger manufacturing sector, a higher share of the economy is isolated from other (dis-)equalizing forces which drive up overall income inequality, but not

[49] A full version showing the coefficients for the year dummies can be found in appendix table 3.A.10.

wage inequality. More surprisingly, trade openness is associated with a higher similarity between income inequality and the Theil index for all measures, and is significant for the SWIID. The absolute effect is rather small, however: A one percentage point increase in the openness ratio implies a 0.1-0.5% lower dissimilarity in the two measures. Nevertheless, it is interesting to note that trade seems to be associated with wage inequality and overall income inequality going in the same direction - although causality could also be the other way around.

The tfp variable, capturing the level of technology, is significantly associated with a smaller gap between wage and income inequality in the SWIID specifications. Its effect is large[50] compared to that of the other variables, and it would appear that technological change is affecting both wage and income inequality in the same way. However, the effect is not robust across data sources with coefficients turning positive in the WIID, LIS, and SILC models.

Table 3.5: Determinants of the difference between wage and income inequality, FE results

	(1) wiid	(2) swiid_n	(3) swiid_m	(4) lis	(5) silc	(6) wb
GDPpc	9.51e-06	1.78e-05	1.77e-05	2.72e-05	2.05e-05	4.40e-05
	(1.03e-05)	(1.25e-05)	(1.24e-05)	(1.97e-05)	(1.40e-05)	(3.07e-05)
Pop. growth	0.192**	0.0493	0.0488	-0.0548	-0.0392	0.0992
	(0.0740)	(0.0503)	(0.0504)	(0.151)	(0.0750)	(0.0724)
Share urban	0.0304*	0.00710	0.00224	-0.0243	-0.0251	0.0200
	(0.0169)	(0.0127)	(0.0119)	(0.0764)	(0.0274)	(0.0153)
Manuf.v.add.	0.00266	0.0173*	0.0193**	0.0307	0.0247	0.0146
	(0.0147)	(0.00964)	(0.00910)	(0.0405)	(0.0260)	(0.0152)
Trade open.	-0.00352	-0.00471**	-0.00481**	-0.00318	-0.00278	-0.00389
	(0.00289)	(0.00234)	(0.00241)	(0.00301)	(0.00255)	(0.00304)
Price lev. inv.	-0.470*	0.485	0.522	0.324	-0.473	-0.0473
	(0.271)	(0.327)	(0.330)	(0.580)	(0.680)	(0.301)
# imputed	-0.0408**	-0.0659***	-0.0659***	0.000215	-0.0355**	-0.0649***
	(0.0172)	(0.0181)	(0.0185)	(0.0136)	(0.0138)	(0.0200)
Tfp	0.895*	-0.513***	-0.460*	0.257	0.349	-0.198
	(0.459)	(0.192)	(0.249)	(1.154)	(0.659)	(0.416)
Constant	4.421***	8.101***	8.490***	4.682	9.531***	8.313***
	(1.173)	(0.737)	(0.718)	(5.502)	(2.312)	(0.982)
Year dummies &other controls	YES	YES	YES	YES	YES	YES
Observations	618	1.521	1.511	120	256	483
R-squared	0,383	0,327	0,323	0,56	0,228	0,358
# of countries	66	82	82	35	28	73

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the logged percentage difference between the (normalized) Theil index and the Gini coefficient from the data source indicated in the top row. Swiid_n and swiid_m refer to net and market inequality, respectively. Silc denotes the EU SILC data and wb the WDI Gini coefficients. The full results, including the year dummies, can be found in 3.A.10.

[50] The variable is for every country normalized to a baseline value of 1 in 2005. A one standard deviation increase, equivalent to 0.27 points on the normalized scale, would lead to a 13.5 percent lower difference between the Theil index and the SWIID indices of income inequality.

Another surprising finding is that the variable measuring the extent of imputation in the underlying data points is negative and significant in all specifications except the LIS. One interpretation of this finding is that, because of the attempt to tamper with the data as little as possible, the Theil index tends to show small changes in inequality in years with more imputations in the sectoral data. Because income inequality is more sluggish than wage inequality, this could mean that the smoother series has a closer resemblance to the developments in income inequality than the more erratic one. Another explanation is that countries where manufacturing wage inequality is closer to overall inequality have more missing data points, which is somewhat puzzling, but fits with the result on sectoral coverage (shown in the random effects results in table 3.A.11). Like in the previous estimations of the relationship between the Theil index and the income inequality measures, the sectoral coverage variable is positive, but significant only for the WIID. Considering the fact that it effectively ranges from 0 to 18, the effect is rather large in the WIID specification (and the LIS, where it is borderline significant). According to the point estimates, the inclusion of one additional sector in the Theil index is associated with a 6 to 7 percent higher difference between wage and income inequality. One would expect that a better sectoral coverage would lead to more accurate numbers of wage inequality and, because wage inequality is a constitutive part of income inequality, this would lead to lower average differences between the two measures.

It is also interesting to note the value of the constant. Focusing on the fixed effects specifications, there appears to be a "baseline" difference between wage and income inequality of roughly between 4.5 and 9.5 percent. Apart from Europe and Central Asia, which display consistently larger differences between the two measures across all specifications, no new insights are obtained about the baseline differences between wage and income inequality from the regional dummies shown in the random effects results (see table 3.A.11). The year dummies do, however, indicate that the differences between the wage inequality index and the income inequality measures are decreasing over time for all measures except the WIID Ginis.

The last thing worth mentioning is that although almost nothing is significant in the LIS specifications, the included variables explain over 50% of the variation in the differences between the two measures. Given that the LIS data are of high quality and the most consistently measured data source of the four measures included, this suggests that the other measures still suffer from a substantial amount of measurement error. Apart from the abovementioned change in the industrial classification scheme to a new and more crude version, this might be the main reason for the lack of a robust association between the Theil index and other measures of income inequality in the previous estimations.

Overall, - measurement error - new classification (given that the new UTIP index also does not)

3.6 Conclusion

The core of this paper is the construction of a Theil index of between-sectoral wage inequality for the manufacturing industry, based on the UNIDO industrial statistics. A very similar index has been built by the University of Texas Inequality Project (UTIP) for the years 1970 to 2008; however, their UTIP-UNIDO index does not include any information on within-sectoral inequality, and it is not clear which sectors are included in the index every year, and in which cases previous versions of the UNIDO industrial statistics have been used directly, or through smoothing out differences with previous versions of the UTIP index. I have therefore recalculated the Theil index for the 48-year time period from 1963 to 2010 for which data was available at the time of writing of this paper. The index relies exclusively on the UNIDO industrial statistics (Rev. 3), which provide data at the 2-digit sector level. I provide detailed information with respect to the sectors covered in each country, as well as the imputation methods used and further variables from the UNIDO incorporated into the same. I then make a recommendation as to which version can be used best in the context of dynamic empirical applications of wage inequality based on an analysis of different versions of the Theil index, reflecting the trade-off between time and sectoral coverage.

The narrow scope of the wage inequality measure, on the one hand, has the advantage of being consistently defined across countries and years, but, on the other hand, restricts the applicability of the index. This paper argues that the latter point is one of the main drawbacks of the index, and presents evidence that its generalizability is severely limited. This applies not only to the extent to which the index allows conjectures about the overall level of income inequality in a society. There is reason to also question the "internal" capability of the index to accurately reflect developments in manufacturing wage inequality. Because it relies on sector-level data on wages and employment which is aggregated at the 2-digit level of industrial classification, the index only measures between-sectoral wage inequality and cannot give account of inequality within sectors. Using data provided at the more disaggregated 3- and 4-digit level, the potential magnitude of within-sectoral inequality is estimated to be at least 40 percent of overall manufacturing wage inequality. Moreover, I find that the between-sectoral index is not generally able to trace changes in between-sectoral inequality over time, contrary to what has been argued by Conceição and Galbraith (2000). Using more detailed sector-level data, I find that looking only at changes in between-sectoral inequality leads to an erroneous image of developments in overall manufacturing wage inequality in around 13% of cases. Given that the sector-level data still do not account for individual inequality within sectors, the true error is likely to be larger and remains open to further exploration.

The analysis of the "external" validity of the index, that is, the extent to which the Theil index is representative of overall income inequality, builds on prior work by Galbraith and Kum (2005), and Galbraith et al. (2015). The authors argue in favor of a stable relationship between the narrowly defined Theil index of wage inequality and the Gini indices of income inequality provided in the WIID and other data sources, which comprise other components besides labor market income. Their finding cannot be confirmed in a broader setting which employs several additional, arguably more consistent, measures of income inequality. Going one step further, this paper tries to find out what causes the measures to show such a weak association with the Theil index, given that there are good theoretical reasons to expect a strong link between them. The deviations between the

Theil index and the other measures of income inequality are regressed on a number of potential explanatory variables. The explanatory power of this model is much lower than for the approach in which both measures are directly regressed on each other, and the coefficient estimates of the explanatory variables are not very insightful. Nevertheless, a few stylized facts emerge. A baseline difference of around 4.5-9.5% between the Theil index and measures of income inequality is found. Furthermore, the signs and robustness of the two "internal" variables, sectoral coverage and the number of imputations, might indicate that there are more missing data points in countries where wage- and income inequality are more similar. Although they do not matter much in the association of the level variables, they are able to explain part of the deviations of the two measures.

In sum, while a measure of between-sectoral wage inequality certainly has its merits and is a valuable resource for empirical analyses with a focus on manufacturing and/or the development of industrial sectors, the general applicability of the index appears to be much more limited than suggested elsewhere. Although all conclusions drawn only apply to the inequality numbers based on the most recent industrial classification scheme used in the UNIDO industrial statistics, given that newly added years after 2010 will - at least for now - be in the new classification scheme as well, the relevance of these results will be growing as the time coverage of the index is extended to more recent years.

3.A Appendix

Table 3.A.1: Deviations between the long and short versions of the Theil index by country

	LONG VERSION							SHORT VERSION								
Country	Max. absolute deviation	Mean absolute deviation	Std. dev.	Corr. with full version	Corr. of diff. with full version	# of years with dev.	Mean Theil index	Max. absolute deviation	Mean absolute deviation	Std. dev.	Corr. with full version	Corr. of diff. with full version	Mean Theil index	# of years with dev.	# of sectors dropped in long version	Recom-mended version for dynamic analysis
CAF	4628.98	1142.2	1489.72	0.05	0.14	8	0.0296	0					0.0604		1	Short
MDA	4546.43	483.7	1005.23	0.22	0.27	25	0.0164	0					0.0919		11	Short
NLD	691.79	81.8	181.68	0.35	0.03	14	0.0103	0					0.0254		5	Short
NZL	649.09	128.9	133.09	0.93	0.96	24	0.0415	0					0.1268		4	Long
BWA	520.4	161.8	130.33	0.96	0.92	13	0.0964	3.54	0.73	1.75	1	1	0.2419	8	11	Long
ROU	402.39	75.6	90.66	0.53	0.92	21	0.0538	0					0.0472		6	Long
HUN	359	27.2	85.61	0.09	0.51	18	0.0246	0					0.0568		5	Short
MUS	235.51	24.4	52.81	0.98	0.96	43	0.0586	0					0.0683		4	Long
MLT	225.96	60.3	62.46	0.68	0.47	14	0.0128	0					0.029		7	Short
MDG	218.44	81.7	68.62	0.96	0.92	29	0.0463	0					0.0427		5	Long
JAM	192.6	56.5	65.51	0.77	0.64	44	0.15	11.89	4.18	4.85	0.99	0.99	0.2936	23	7	Short
SWE	102.18	16.2	23.8	0.91	0.41	20	0.0054	0					0.0083		5	Short
FIN	97.21	21.9	31.3	0.76	0.8	19	0.0107	0.21	0.03	0.11	1	1	0.0125	15	5	Short
MEX	96.72	17.2	28.82	0.96	0.87	22	0.0517	0.04	0.04				0.0641	1	5	Long
MNG	76.36	13.4	21.2	0.91	0.9	17	0.0842	2.72	-0.36	0.82	1	1	0.09	16	5	Long
DZA	60.09	20.1	25.27	0.83	0.78	14	0.0165	0					0.008		10	Short
LSO	59.87	47.8	9.72	0.81	0.49	9	0.2062	0					0.1371		5	Short
MWI	57.98	15.2	16.39	1	0.98	32	0.0948	6.09	0.59	1.76	1	1	0.0812	25	4	Long
AZE	51.08	18	14.36	0.92	0.86	16	0.1113	0					0.1608		5	Long
AUS	49.72	14.1	19.94	1	1	14	0.0526	0					0.1456		4	Long
FRA	46.83	24	10.49	0.86	0.8	20	0.0285	0					0.0177		5	Short
SWZ	43.11	9.4	14.72	0.97	0.94	11	0.1115	0					0.1003		7	Long
BRA	40.89	8	9.4	0.7	0.39	15	0.124	0					0.122		5	Short
GBR	40.48	9.6	9.82	0.86	0.83	17	0.0144	0					0.0182		5	Short
SGP	40.35	32.8	4.08	0.95	0.66	20	0.0598	2.7	-1.53	0.68	1	1	0.0374	20	5	Short
AUT	37.14	14.5	11.25	0.77	0.77	20	0.0166	0.94	-0.07	0.4	1	1	0.0197	15	5	Short
ALB	29.94	7.5	12.23	0.98	0.98	11	0.0689	0.65	0.06	0.61	1	1	0.1045	3	9	Long
SVK	28.02	11.7	8.05	0.98	0.84	17	0.0249	0					0.0277		1	Short
HRV	27.2	8.6	6.92	0.95	0.89	14	0.0296	0					0.0406		5	Long
MYS	26.59	19.8	3.27	0.89	0.77	11	0.0332	0					0.0342		5	Short
POL	26.44	6.2	8.71	0.96	0.87	18	0.0154	0					0.0282		5	Long
CHN	24.94	13.6	9.95	0.84	0.99	8	0.0785	0					0.0292		5	Long
HTI*	24.35	10	11.88	0.94	0.99	10	0.104	3.48	2.04	0.6	1	1	0.1096	10	2	Long
ESP	23.75	11	4.91	0.98	0.92	17	0.0276	0					0.0277		5	Long

UGA*	22.83	6.4	8.48	1	1	23	0.109						0.1092		4	Long
KOR	20.83	11	5.43	0.98	0.87	17	0.024	0					0.0212		5	Long
SUR*	19.11	3.5	5.97	1	0.98	19	0.0488						0.0514		7	Long
SEN*	18.46	7.7	7.26	0.92	0.97	5	0.041						0.0418		5	Long
ISR	17.23	4	5.2	0.99	0.92	46	0.0485	0					0.0713		4	Long
TTO*	17.14	2.7	5.86	1	0.99	8	0.1525						0.1569		5	Long
TON*	16.33	4.9	6.66	0.99	0.98	9	0.0626						0.0627		5	Long
MAC	16.27	6.1	5.31	1	0.99	11	0.014	0					0.0258		2	Long
PRT	15.37	10.3	2.67	0.99	0.98	15	0.0489	0					0.0508		5	Long
GRC	14.14	6.1	3.46	0.98	0.93	15	0.0284	0.07	-0.03	0.05	1	1	0.031	6	5	Long
PHL	14.08	9.1	5.21	0.99	0.98	13	0.055	0					0.0503		5	Long
SVN	13.63	2	3.59	1	0.98	24	0.0246	2.41	-1.08	0.79	1	1	0.0275	11	6	Long
ITA	13.6	6.2	4.2	0.93	0.98	19	0.0179	0					0.0186		5	Long
NOR	13.55	6	7.27	0.99	0.99	18	0.0107	0					0.014		5	Long
ECU	13.06	5	4.36	0.99	1	14	0.0408	0					0.0547		5	Long
LUX	12.45	4	4.98	1	1	20	0.03	0					0.039		4	Long
IDN	11.35	4	3.94	1	1	20	0.0854	0					0.0968		6	Long
BOL	11.3	3.8	5.55	0.92	0.94	7	0.0543	0					0.0789		5	Long
TUN	11.14	4.9	4	1	1	33	0.1527	0					0.1276		4	Long
IRL*	10.61	5	4.6	0.98	0.91	19	0.0166	10.48	-4.02	3.37	0.99	0.97	0.0154	19	5	Long
COUNTRIES WITH ABSOLUTE DEVIATIONS OF LESS THAN 10%																
USA	9.11	5.3	4.97	0.98	0.95	11	0.0251								5	
JPN	8.76	3.6	3.15	1	0.99	26	0.0418								6	
TUR	8.6	3.2	2.42	1	1	18	0.0505								5	
MKD	8.37	2.8	3.73	0.99	0.97	11	0.0634								5	
JOR	8.1	3	3.21	0.99	1	37	0.0839								6	
UKR	8.08	5.8	2	1	0.99	11	0.0465								5	
DNK	7.97	2.9	3.64	0.99	0.99	18	0.007								5	
EGY	7.82	2.1	2.91	1	0.99	14	0.0523								5	
LBY	7.6	7.6				1	0.0376								1	
THA	7	4.9	2.8	1	0.98	11	0.0578								5	
PNG	6.97	2.8	2.19	1	1	15	0.0806								2	
MAR	6.89	3.6	2.25	1	0.98	12	0.0911								6	
COL	6.71	4.2	1.35	1	0.99	11	0.0346								4	
LKA	6.44	4.2	4.12	1	1	5	0.062								5	
PSE	6.36	4	3.65	1	0.99	14	0.0377								6	
ZAF	6.34	3.5	1.57	1	0.99	20	0.0566								5	
ARG	6.24	2.9	1.47	1	1	10	0.0524								4	
PER	6.18	4	1.75	1	0.99	12	0.2091								5	
BRB	5.57	2.3	2.52	0.99	0.99	11	0.055								5	
CZE	5.05	1.1	1.28	1	0.99	20	0.0099								6	
URY	5.02	3.1	1.4	1	1	11	0.0452								5	
CAN	4.97	2.1	1.88	0.99	0.99	21	0.0183								4	
PAN	4.9	1.3	1.82	1	1	14	0.0589								2	
SYR	4.87	2	2.28	1	1	15	0.1205								4	
TZA	4.81	3.5	0.98	1	1	5	0.0803								2	

BGD	4.76	3.5	3.42	1	1	4	0.0299	5
BGR	4.68	1.6	1.97	1	1	16	0.0841	6
KGZ	4.62	1.3	1.73	1	1	17	0.1636	5
KWT	4.42	3.6	0.74	1	0.89	3	0.2844	5
IND	4.33	1.4	1.94	0.98	0.95	12	0.0778	5
NPL	4.04	1	1.89	1	1	6	0.0635	9
CYP	3.97	1.5	2	1	1	12	0.0286	5
IRN	3.54	1.6	0.86	1	1	12	0.035	5
BEL	3.28	1.1	1.38	1	1	15	0.0566	5
FJI	2.87	0.9	1.06	1	1	32	0.5819	5
SOM	3.16	3.2				1	0.044	2
CHE	2.83	1.7	1.11	1	1	3	0.0229	1
CIV	2.58	1.9	0.7	1	1	4	0.052	2
QAT	2.06	0.9	0.62	1	1	11	0.3886	7
GEO	1.94	0.6	0.85	1	1	5	0.061	2
PAK	1.64	1.6				1	0.0556	4
CHL	1.55	1.2	0.71	1	1	8	0.0606	4
RUS	1.44	0.6	0.45	1	1	9	0.0494	4
PRY	0.72	0.7				1	0.0277	3

Notes.

HTI: Sectors 16, 18, 28, 29, and 34 have imputed values for years 1988-1997. A short version of the index (including sectors 19 and 32) would therefore mean that those other sectors should be dropped, and consistency would still not be established. One alternative would be to have two different, non-comparable short versions for Haiti: one from 1988-1997, and one from 1968-1987. The long version of the index displays larger deviations of around 20% only in the last 4 years (1994-1997), despite the fact that the same 2 sectors are omitted throughout - sectoral composition is therefore not driving the differences in the deviations between the long- and full versions. Because of the high correlation between the two measures over time in both levels and differences, retaining the long version seems justifiable.

UGA: For Uganda, the long version is retained despite deviations of up to 22% in first few years since including the sector causing this deviation (27) would effectively mean a time coverage of only four years from 1963-1966. Additionally, the contribution of the sector causing the deviation is vanishing over time and in the second time spell where data for the sector is present, the deviations are very small (between 2 and 7.5%), in line with the decrease in absolute size of the sector. Additionally, Uganda has a low average deviation of around 6% and almost perfect co-movement of the long- with the full version over time, as indicated by correlation coefficients which round up to one at 2 digits.

SUR: In Suriname, the index would decrease from 20 to only four years of data coverage in the short version. I have therefore decided to keep the long version given that the maximum deviation of 19% only arises in the first year of data (1974) and keeps decreasing thereafter to around 13% in 1975 and 1976 and 10% in 1977. Assuming that the downward trend continues, sacrificing 16 years of data for achieving higher accuracy of supposedly less than 10% seems unreasonable. The very high correlation of the long- with the full version in both levels and differences also supports the long version.

SEN: In Senegal, keeping the sectors causing the deviation of around 18.5% would leave only 5 years of data (1998-2002). Given that only a single year has such a high deviation (again, this is not because more sectors are omitted in that year) and the correlations over time are fairly high, the long version is retained.

TTO: In Trinidad and Tobago, only a single year (1998) is causing the deviation of around 17%. Upon closer inspection of the data, this deviation can be traced back to what is likely to be a glitch in the data, with employee numbers in sector 35 suddenly dropping to 16 (160 being a much more reasonable number) before rising again to 176 in 1999. This drop is also not warranted by changes in any other variables, or by a similar drop in other sectors in that year.

TON: In Tonga, sector 35 is responsible for the one-year deviation of around 16% in 1991. The contribution of the sectors is decreasing thereafter and the deviations are very small. While this does not point towards a lower contribution of the sector in the years preceding 1991, keeping only the years 1991-2004 for which data are provided in sector 35 would lead to another problem: many other sectors have 0s for wages and employees in the later years, making the short index not very informative for the overall level and development of inequality in the country. Given the low average deviation of less than 5%, and the high correlation of the long- and full indices, retaining the long version therefore seems like the better option.

IRL: In Ireland, changing to the short version requires the dropping of two sectors (23 and 36) which have a lot of imputed data in the years covered by the "short" version. Omitting the dropped sectors would result in a deviation of approximately the same magnitude (10.5%), but fewer years (18 instead of 46).

Table 3.A.2: Overview of the Theil index by country

country code	country	years	sectors	imputed	dynamic version	region	dev. status	Theil(normalized)	Theil	standard deviation Theil
AFG	Afghanistan	9	9	6	long	SA	developing	0.0023	0.00497	0.00519
ALB	Albania	18	12	40	long	ECANA	developing	0.0292	0.06891	0.05679
ARG	Argentina	19	18	72	long	LAC	developed	0.0182	0.05239	0.01377
AUS	Australia	44	18	184	long	EAP	developed	0.0181	0.05263	0.08744
AUT	Austria	20	18	80	short	ECANA	developed	0.0067	0.01973	0.00277
AZE	Azerbaijan	21	18	49	long	ECANA	developing	0.0386	0.11126	0.0596
BDI	Burundi	23	18	166	long	SSA	developing	0.0235	0.06206	0.0324
BEL	Belgium	47	18	130	long	ECANA	developed	0.0181	0.05659	0.06289
BEN	Benin	8	18	9	long	SSA	developing	0.0256	0.07825	0.01784
BFA	Burkina Faso	10	18	0	long	SSA	developing	0.0115	0.03322	0.02159
BGD	Bangladesh	32	18	144	long	SA	developing	0.0096	0.02987	0.02131
BGR	Bulgaria	48	17	55	long	ECANA	developing	0.0293	0.08411	0.04892
BLZ	Belize	4	16	24	long	LAC	developing	0.035	0.1149	0.07713
BOL	Bolivia	32	18	22	long	LAC	developing	0.0188	0.05431	0.02979
BRA	Brazil	15	18	636	short	LAC	developing	0.0394	0.12196	0.01522
BRB	Barbados	28	12	0	long	LAC	developed	0.0221	0.055	0.01579
BWA	Botswana	30	8	73	long	SSA	developing	0.038	0.0964	0.08864
CAF	Central African Republic	8	16	64	short	SSA	developing	0.0211	0.06039	0.01242
CAN	Canada	48	18	24	long	ECANA	developed	0.0063	0.01831	0.0037
CHE	Switzerland	11	19	158	long	ECANA	developed	0.0034	0.02285	0.01407
CHL	Chile	46	18	65	long	LAC	developed	0.021	0.06057	0.02375
CHN	China	34	18	294	long	EAP	developing	0.0272	0.07853	0.09696
CIV	Cte d'Ivoire	22	15	9	long	SSA	developing	0.0194	0.05199	0.01735
CMR	Cameroon	33	18	205	long	SSA	developing	0.0382	0.10795	0.06666
COG	Congo	21	14	142	long	SSA	developing	0.0222	0.06684	0.02836
COL	Colombia	48	18	78	long	LAC	developing	0.0119	0.03459	0.00721
CRI	Costa Rica	41	18	364	long	LAC	developing	0.0118	0.05152	0.02943
CUB	Cuba	15	14	42	long	LAC	developing	0.0015	0.00477	0.00293
CYP	Cyprus	48	18	30	long	ECANA	developed	0.0099	0.02861	0.00953
CZE	Czech Republic	21	17	38	long	ECANA	developed	0.0035	0.00988	0.00389
DEU	Germany	27	18	72	long	ECANA	developed	0.0015	0.00438	0.00563
DNK	Denmark	47	18	152	long	ECANA	developed	0.0022	0.00699	0.00289
DOM	Dominican Rep.	23	18	0	long	LAC	developing	0.0219	0.06321	0.02334
DZA	Algeria	14	8	48	short	MENA	developing	0.0028	0.00802	0.00273
ECU	Ecuador	46	18	2	long	LAC	developing	0.0141	0.04084	0.01833
EGY	Egypt	47	18	256	long	MENA	developing	0.0147	0.05229	0.04363
ERI	Eritrea	19	23	9	long	SSA	developing	0.0134	0.04191	0.02442
ESP	Spain	47	18	6	long	ECANA	developed	0.0095	0.02757	0.00788
EST	Estonia	19	22	337	long	ECANA	developed	0.0411	0.13342	0.16341

country code	country	years	sectors	imputed	dynamic version	region	dev. status	Theil(normalized)	Theil	standard deviation Theil
ETH	Ethiopia	20	23	7	long	SSA	developing	0.0117	0.0365	0.02189
FIN	Finland	19	18	24	short	ECANA	developed	0.0041	0.01247	0.00482
FJI	Fiji	42	15	276	long	EAP	developing	0.0203	0.05819	0.04232
FRA	France	20	18	339	short	ECANA	developed	0.006	0.01766	0.00208
GAB	Gabon	16	18	181	long	SSA	developing	0.0326	0.10511	0.04791
GBR	United Kingdom	17	18	171	short	ECANA	developed	0.0059	0.01825	0.00366
GEO	Georgia	13	21	20	long	ECANA	developing	0.0202	0.06097	0.0227
GHA	Ghana	25	18	0	long	SSA	developing	0.0328	0.09483	0.03213
GMB	Gambia	8	18	0	long	SSA	developing	0.0049	0.01419	0.00389
GRC	Greece	45	18	146	long	ECANA	developed	0.0098	0.02839	0.00401
GTM	Guatemala	31	18	180	long	LAC	developing	0.0284	0.07959	0.06715
HKG	Hong Kong	35	18	29	long	EAP	developed	0.0092	0.02642	0.03769
HND	Honduras	34	18	233	long	LAC	developing	0.0239	0.06183	0.03547
HRV	Croatia	25	18	12	long	ECANA	developed	0.0103	0.02959	0.01218
HTI	Haiti	30	17	121	long	LAC	developing	0.0415	0.10405	0.09331
HUN	Hungary	18	18	112	short	ECANA	developed	0.0168	0.05677	0.02448
IDN	Indonesia	40	17	57	long	EAP	developing	0.0305	0.08543	0.03494
IND	India	47	18	8	long	SA	developing	0.027	0.07777	0.01954
IRL	Ireland	47	18	77	long	ECANA	developed	0.0058	0.01658	0.00319
IRN	Iran	43	18	36	long	MENA	developing	0.012	0.03503	0.0227
IRQ	Iraq	30	18	108	long	MENA	developing	0.0082	0.02301	0.01413
ISL	Iceland	29	16	39	long	ECANA	developed	0.0086	0.02371	0.01183
ISR	Israel	47	16	36	long	MENA	developed	0.0173	0.04853	0.02216
ITA	Italy	43	18	18	long	ECANA	developed	0.0062	0.01793	0.00581
JAM	Jamaica	34	11	233	short	LAC	developing	0.1169	0.29362	0.15091
JOR	Jordan	48	17	87	long	MENA	developing	0.0303	0.08387	0.02922
JPN	Japan	48	17	12	long	EAP	developed	0.0148	0.04183	0.02265
KAZ	Kazakhstan	10	23	0	long	ECANA	developing	0.021	0.0657	0.02997
KEN	Kenya	48	18	400	long	SSA	developing	0.0234	0.0787	0.02738
KGZ	Kyrgyzstan	21	18	58	long	ECANA	developing	0.0569	0.16359	0.18953
KOR	Korea	44	18	0	long	EAP	developed	0.0083	0.02396	0.00564
KWT	Kuwait	44	17	210	long	MENA	developed	0.0912	0.28442	0.14956
LBR	Liberia	3	18	0	long	SSA	developing	0.0192	0.0554	0.01436
LBY	Libya	17	16	32	long	MENA	developing	0.0141	0.03762	0.03062
LKA	Sri Lanka	41	18	244	long	SA	developing	0.0215	0.06201	0.02747
LSO	Lesotho	9	7	27	short	SSA	developing	0.0566	0.13708	0.04618
LTU	Lithuania	19	23	66	long	ECANA	developed	0.014	0.04312	0.01508
LUX	Luxembourg	25	13	52	long	ECANA	developed	0.0141	0.02996	0.0307
LVA	Latvia	19	21	101	long	ECANA	developed	0.0139	0.04154	0.04273
MAC	Macao	30	18	94	long	EAP	developed	0.0056	0.01396	0.01243
MAR	Morocco	35	17	64	long	MENA	developing	0.0325	0.09112	0.03225

country code	country	years	sectors	imputed	dynamic version	region	dev. status	Theil(normalized)	Theil	standard deviation Theil
MDA	Moldova	10	3	18	short	ECANA	developing	0.036	0.09187	0.05026
MDG	Madagascar	29	15	75	long	SSA	developing	0.012	0.04632	0.05341
MEX	Mexico	27	18	224	long	LAC	developing	0.0146	0.05173	0.03196
MKD	Macedonia	21	18	97	long	ECANA	developing	0.022	0.06341	0.03355
MLT	Malta	14	16	54	short	MENA	developed	0.0094	0.029	0.02796
MNG	Mongolia	19	18	102	long	EAP	developing	0.0312	0.08422	0.03902
MOZ	Mozambique	26	18	252	long	SSA	developing	0.0507	0.1965	0.19345
MUS	Mauritius	43	15	20	long	SSA	developing	0.0218	0.05858	0.0268
MWI	Malawi	40	11	58	long	SSA	developing	0.0403	0.09476	0.06529
MYS	Malaysia	11	18	46	short	EAP	developing	0.011	0.03418	0.00308
NGA	Nigeria	34	18	216	long	SSA	developing	0.0106	0.02885	0.01468
NIC	Nicaragua	21	18	0	long	LAC	developing	0.005	0.01453	0.00502
NLD	Netherlands	14	18	66	short	ECANA	developed	0.0087	0.02539	0.02359
NOR	Norway	46	18	78	long	ECANA	developed	0.0037	0.01067	0.006
NPL	Nepal	13	14	58	long	SA	developing	0.0254	0.06354	0.03242
NZL	New Zealand	47	18	295	long	EAP	developed	0.0165	0.0415	0.08545
OMN	Oman	18	22	22	long	MENA	developed	0.0357	0.10944	0.03435
PAK	Pakistan	44	18	432	long	SA	developing	0.0159	0.05563	0.02346
PAN	Panama	43	18	228	long	LAC	developing	0.0217	0.05888	0.02269
PER	Peru	28	18	177	long	LAC	developing	0.0676	0.20908	0.16057
PHL	Philippines	46	18	152	long	EAP	developing	0.0188	0.05496	0.01344
PNG	Papua New Guinea	25	16	0	long	EAP	developing	0.0291	0.08062	0.02301
POL	Poland	40	18	130	long	ECANA	developed	0.005	0.01541	0.01157
PRI	Puerto Rico	20	18	144	long	LAC	developed	0.0319	0.1185	0.08301
PRT	Portugal	20	18	142	long	ECANA	developed	0.0172	0.04886	0.01008
PRY	Paraguay	2	17	0	long	LAC	developing	0.0098	0.02766	0.00007
QAT	Qatar	25	13	189	long	MENA	developed	0.0134	0.38861	0.05925
ROU	Romania	33	17	124	long	ECANA	developing	0.1512	0.0538	0.07169
RUS	Russia	15	18	0	long	ECANA	developed	0.0191	0.04944	0.01314
SEN	Senegal	29	18	120	long	SSA	developing	0.0171	0.04101	0.02397
SGP	Singapore	20	18	40	short	EAP	developed	0.0158	0.03737	0.00509
SLV	El Salvador	36	18	216	long	LAC	developing	0.0124	0.0407	0.02331
SOM	Somalia	14	16	6	long	SSA	developing	0.0141	0.04401	0.02047
SRB	Serbia and Montenegro	12	18	90	long	ECANA	developing	0.016	0.12232	0.11691
SUR	Suriname	20	11	0	long	LAC	developing	0.0427	0.04876	0.02528
SVK	Slovakia	17	21	18	short	ECANA	developed	0.0203	0.02772	0.00936
SVN	Slovenia	24	17	66	long	ECANA	developed	0.0091	0.02459	0.00939
SWE	Sweden	20	18	38	short	ECANA	developed	0.009	0.00829	0.00549
SWZ	Swaziland	24	5	30	long	SSA	developing	0.0027	0.11148	0.04479
SYR	Syria	48	18	267	long	MENA	developing	0.069	0.12055	0.05836
THA	Thailand	39	18	588	long	EAP	developing	0.0458	0.05783	0.02888

country code	country	years	sectors	imputed	dynamic version	region	dev. status	Theil(normalized)	Theil	standard deviation Theil
TON	Tonga	30	18	175	long	EAP	developing	0.0225	0.06259	0.06513
TTO	Trinidad and Tobago	42	17	257	long	LAC	developed	0.0223	0.15252	0.08948
TUN	Tunisia	37	14	116	long	MENA	developing	0.0505	0.15268	0.13522
TUR	Turkey	47	18	42	long	ECANA	developing	0.0369	0.05052	0.03092
TWN	Taiwan	29	18	72	long	EAP	developed	0.0173	0.01477	0.00341
TZA	Tanzania	43	18	349	long	SSA	developing	0.0051	0.08027	0.03656
UGA	Uganda	23	10	26	long	SSA	developing	0.0279	0.10898	0.0796
UKR	Ukraine	19	18	6	long	ECANA	developing	0.0488	0.04651	0.01613
URY	Uruguay	41	18	216	long	LAC	developed	0.0161	0.04516	0.01753
USA	United States	45	18	54	long	ECANA	developed	0.0158	0.02505	0.00448
VEN	Venezuela	35	18	72	long	LAC	developed	0.0086	0.04673	0.01972
YEM	Yemen	9	17	18	long	MENA	developing	0.0156	0.08202	0.02211
ZAF	South Africa	48	18	223	long	SSA	developing	0.0289	0.05658	0.00942
ZMB	Zambia	22	17	146	long	SSA	developing	0.0046	0.0518	0.01617

Notes. The column "years" is the number of total (not necessarily consecutive) years covered for each country. "Sectors" refers to the number of ISIC 2-digit level sectors on which the Theil index is based. "Imputed" contains the total number of imputed data points across all sectors and years. It should be noted that this number tends rise with higher time coverage. "dynv" is short for "dynamic version" and contains the recommendation as to which version in the case of deviations between the two version of more than 10% in any year, with the exceptions discussed in appendix table 3.A.1. "Region" refers to the geographic region and relies on the World Bank classification. SSH=Sub-Saharan Africa, SA=South Asia, LAC=Latin America and Caribbean, ECANA=Europe, Central Asia, and North America, MENA=Middle East and North Africa, and EAP=East Asia and Pacific. "Devstat" refers to the classification of countries as developed or developing and relies on the World Bank categorization, which is based on GNI. Theil(n) is the normalized version of the Theil index. The standard deviation in the last column is for the non-normalized version of the Theil.

Table 3.A.3: Correlation with the UTIP index and extent of imputation

country code	# of imputed data points	# of years	# of years in UTIP	correlation with UTIP, levels	correlation with UTIP, differences	mean %-devations from UTIP, levels	mean %-devations from UTIP, differences
AFG	6	9	22	0.982	0.979	59.3	59.3
ALB	40	18	19	0.973	0.974	25.9	25.9
ARG	72	19	17	0.999	0.997	0.9	0.9
AUS	184	44	40	0.568	-0.393	21.1	21.1
AUT	80	20	44	0.673	0.584	12.5	12.5
AZE	49	21	17	0.973	0.749	11.3	11.3
BDI	166	23	17	1	1	0	0
BEL	130	47	42	0.136	-0.054	22.3	22.3
BEN	9	8	7	1	1	0	0
BFA	0	10	10	1	1	0	0
BGD	144	32	28	0.997	0.97	0.6	0.6
BGR	55	48	45	-0.169	-0.152	47.5	47.5
BLZ	24	4	2	1	n/a	0	0
BOL	22	32	32	0.882	0.366	24.4	24.4
BRA	636	15	17	0.879	0.919	13.1	13.1
BRB	0	28	28	0.984	0.988	3.4	3.4
BWA	73	30	27	0.048	0.234	62.5	62.5
CAF	64	8	19	0.984	0.997	28.9	28.9
CAN	24	48	45	0.972	0.906	2	2
CHE	158	11	5	0.872	0.954	3.5	3.5
CHL	65	46	44	0.994	0.988	5.7	5.7
CHN	294	34	16	0.997	0.61	2.4	2.4
CIV	9	22	22	1	0.999	0.3	0.3
CMR	205	33	28	1	0.998	1.4	1.4
COG	142	21	14	1	1	0	0
COL	78	48	43	0.997	0.996	0.6	0.6
CRI	364	41	22	1	1	0.1	0.1
CUB	42	15	13	1	1	0	0
CYP	30	48	46	0.948	0.825	19.8	19.8
CZE	38	21	20	0.991	0.831	10	10
DEU	72	27	30	-0.454	-0.025	344.6	344.6
DNK	152	47	42	0.998	0.998	3	3
DOM	0	23	23	1	1	0	0
DZA	48	14	27	0.997	0.998	2.7	2.7
ECU	2	46	45	0.997	0.997	1.3	1.3
EGY	256	47	39	1	0.997	0.2	0.2
ERI	9	19	42	0.734	0.966	26.3	26.3
ESP	6	47	45	0.992	0.994	2.7	2.7
EST	337	19	9	-0.027	-0.599	29.3	29.3
ETH	7	20	44	0.996	0.995	0.4	0.4
FIN	24	19	45	0.931	0.751	7.5	7.5
FJI	276	42	32	0.717	0.861	13.7	13.7
FRA	339	20	30	0.902	0.934	11.8	11.8
GAB	181	16	8	1	1	0	0
GBR	171	17	41	0.647	0.508	12.7	12.7
GEO	20	13	11	0.998	0.999	1.6	1.6
GHA	0	25	28	0.191	0.619	35.6	35.6
GMB	0	8	8	0.994	1	4.2	4.2
GRC	146	45	41	0.987	0.979	1.1	1.1
GTM	180	31	26	0.996	0.831	24.1	24.1
HKG	29	35	36	0.847	0.692	12.9	12.9

country code	# of imputed data points	# of years	# of years in UTIP	correlation with UTIP, levels	correlation with UTIP, differences	mean %-devations from UTIP, levels	mean %-devations from UTIP, differences
HND	233	34	26	0.178	0.1	41.1	41.1
HRV	12	25	23	0.996	0.984	3.6	3.6
HTI	121	30	21	0.288	0.423	23.8	23.8
HUN	112	18	43	0.954	0.694	8.2	8.2
IDN	57	40	36	0.903	0.922	3.8	3.8
IND	8	47	45	0.999	0.988	0.3	0.3
IRL	77	47	45	0.858	0.921	12.3	12.3
IRN	36	43	42	0.999	0.998	1.4	1.4
IRQ	108	30	27	1	1	0	0
ISL	39	29	20	0.985	0.941	3.4	3.4
ISR	36	47	44	0.996	0.944	2.1	2.1
ITA	18	43	40	0.997	0.989	1.7	1.7
JAM	233	34	34	-0.136	-0.032	57.5	57.5
JOR	87	48	42	0.743	0.471	8.1	8.1
JPN	12	48	45	0.999	0.998	1.8	1.8
KAZ	0	10	10	1	1	0	0
KEN	400	48	40	0.602	0.879	6.1	6.1
KGZ	58	21	13	0.856	0.961	24.2	24.2
KOR	0	44	44	0.995	0.989	2.1	2.1
KWT	210	44	35	1	1	0	0
LBR	0	3	n/a	n/a	n/a	n/a	n/a
LBY	32	17	17	0.984	0.94	6.7	6.7
LKA	244	41	26	0.999	1	0.5	0.5
LSO	27	9	14	0.973	0.964	4	4
LTU	66	19	16	0.984	0.954	1.5	1.5
LUX	52	25	45	0.105	0.288	20.8	20.8
LVA	101	19	16	0.686	0.143	22.8	22.8
MAC	94	30	26	0.971	0.952	13.7	13.7
MAR	64	35	33	0.99	0.961	4.7	4.7
MDA	18	10	17	0.556	-1	20.8	20.8
MDG	75	29	26	0.661	0.784	64.3	64.3
MEX	224	27	31	0.887	0.249	9.6	9.6
MKD	97	21	20	0.986	0.959	9.9	9.9
MLT	54	14	44	0.979	0.991	19.3	19.3
MNG	102	19	17	0.928	0.907	8.5	8.5
MOZ	252	26	13	0.966	1	2.1	2.1
MUS	20	43	40	0.992	0.987	9.6	9.6
MWI	58	40	35	0.971	0.958	15.1	15.1
MYS	46	11	39	0.994	0.996	2.5	2.5
NGA	216	34	28	1	1	0	0
NIC	0	21	21	1	1	0	0
NLD	66	14	43	0.473	-0.129	43.4	43.4
NOR	78	46	44	0.304	0.057	6.4	6.4
NPL	58	13	10	1	1	0.3	0.3
NZL	295	47	41	0.467	0.457	24.6	24.6
OMN	22	18	15	1	1	0	0
PAK	432	44	32	1	1	0.1	0.1
PAN	228	43	40	0.864	0.848	36	36
PER	177	28	21	1	0.996	1.6	1.6
PHL	152	46	41	0.964	0.997	2.4	2.4
PNG	0	25	27	0.997	0.998	1.7	1.7
POL	130	40	37	0.993	0.902	2.7	2.7

country code	# of imputed data points	# of years	# of years in UTIP	correlation with UTIP, levels	correlation with UTIP, differences	mean %-devations from UTIP, levels	mean %-devations from UTIP, differences
PRI	144	20	12	1	1	0	0
PRT	142	20	27	0.089	-0.686	7.3	7.3
PRY	0	2	3	-1		13.7	13.7
PSE	9	14	15	0.987	0.981	8.3	8.3
QAT	189	25	15	1	1	0.4	0.4
ROU	124	33	26	0.942	0.699	14.6	14.6
RUS	0	15	44	0.999	0.999	0.8	0.8
SEN	120	29	29	0.979	0.968	17.2	17.2
SGP	40	20	46	0.986	0.968	6.1	6.1
SLV	216	36	28	1	1	0	0
SOM	6	14	12	0.978	0.966	2	2
SRB	90	12	n/a	n/a	n/a	n/a	n/a
SUR	0	20	24	0.997	0.985	3.3	3.3
SVK	18	17	17	0.918	0.823	14.6	14.6
SVN	66	24	22	0.933	0.86	13.6	13.6
SWE	38	20	38	0.903	0.67	21.1	21.1
SWZ	30	24	26	0.985	0.968	4.9	4.9
SYR	267	48	28	0.94	0.719	43.8	43.8
THA	588	39	23	0.885	0.622	7	7
TON	175	30	23	0.998	0.996	2.5	2.5
TTO	257	42	26	0.997	0.992	0.7	0.7
TUN	116	37	29	1	0.999	6.2	6.2
TUR	42	47	43	1	0.998	0.5	0.5
TWN	72	29	25	1	1	0	0
TZA	349	43	34	0.875	0.878	9.1	9.1
UGA	26	23	21	-0.016	-0.01	56.1	56.1
UKR	6	19	19	0.992	0.996	4.1	4.1
URY	216	41	32	0.985	0.859	2.9	2.9
USA	54	45	42	0.714	0.424	3	3
VEN	72	35	34	0.938	0.837	38.1	38.1
YEM	18	9	10	0.051	0.431	33.5	33.5
YUG	0	27	35	1	1	0	0
ZAF	223	48	41	0.969	0.953	1.8	1.8
ZMB	146	22	18	0.983	0.977	7.4	7.4

Table 3.A.4: Imputed values vs. dropping of sectors: RE and FE results

	(1)		(2)		(3)		(4)	
	RE	se	FE	se, r	RE	se	RE	se, r
Imputations	3.030***	(0.466)	3.030**	(1.225)	3.111***	(0.460)	3.111***	(1.154)
Dropped sectors	-5.282***	(0.815)	-5.282	(3.444)	-2.396***	(0.607)	-2.396	(1.817)
1964.year	1.090	(13.44)	1.090	(1.367)				
1965.year	0.915	(13.25)	0.915	(1.908)				
1966.year	-0.912	(13.15)	-0.912	(2.741)				
1967.year	2.733	(12.86)	2.733	(3.271)				
1968.year	-0.793	(12.36)	-0.793	(3.941)				
1969.year	-1.377	(12.61)	-1.377	(4.602)				
1970.year	-0.557	(12.48)	-0.557	(4.956)				
1971.year	-2.374	(12.34)	-2.374	(5.034)				
1972.year	0.193	(12.41)	0.193	(5.308)				
1973.year	3.219	(12.34)	3.219	(6.548)				
1974.year	3.266	(12.22)	3.266	(6.888)				
1975.year	2.218	(12.20)	2.218	(7.074)				
1976.year	3.129	(12.21)	3.129	(7.159)				
1977.year	0.184	(12.13)	0.184	(7.882)				
1978.year	4.355	(12.13)	4.355	(7.570)				
1979.year	-0.618	(12.13)	-0.618	(6.960)				
1980.year	4.184	(12.07)	4.184	(8.647)				
1981.year	2.055	(12.07)	2.055	(9.078)				
1982.year	0.204	(12.10)	0.204	(9.514)				
1983.year	-4.433	(12.10)	-4.433	(9.099)				
1984.year	-5.342	(12.10)	-5.342	(10.06)				
1985.year	-41.32***	(12.09)	-41.32	(37.52)				
1986.year	-6.042	(12.08)	-6.042	(13.27)				
1987.year	-7.133	(12.05)	-7.133	(13.97)				
1988.year	-13.52	(12.20)	-13.52	(16.13)				
1989.year	-8.212	(12.17)	-8.212	(14.38)				
1990.year	1.174	(12.05)	1.174	(7.983)				
1991.year	2.955	(11.99)	2.955	(8.181)				
1992.year	2.617	(12.13)	2.617	(8.634)				
1993.year	5.108	(12.10)	5.108	(9.354)				
1994.year	-4.776	(11.97)	-4.776	(9.583)				
1995.year	-6.343	(12.20)	-6.343	(11.63)				
1996.year	11.93	(12.19)	11.93	(13.04)				
1997.year	12.82	(12.26)	12.82	(14.22)				
1998.year	18.33	(12.29)	18.33	(15.03)				
1999.year	17.17	(12.58)	17.17	(15.78)				
2000.year	19.42	(12.39)	19.42	(15.90)				
2001.year	18.17	(12.49)	18.17	(15.88)				
2002.year	22.11*	(12.69)	22.11	(16.88)				
2003.year	17.89	(12.79)	17.89	(17.33)				
2004.year	22.07*	(12.80)	22.07	(16.96)				
2005.year	18.88	(12.84)	18.88	(17.25)				
2006.year	21.41*	(12.95)	21.41	(17.46)				
2007.year	23.76*	(13.45)	23.76	(18.16)				
2008.year	5.556	(17.10)	5.556	(25.21)				
Constant	0.0131	(10.30)	0.0131	(5.216)	0.105	(4.026)	0.105	(4.081)
Observations	3,627		3,627		3,627		3,627	
# of countries	135		135		135		135	
R^2 (within)	0.036		0.036		0.016		0.016	

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the percentage deviation between the dynamic version of the newly constructed Theil index and the UTIP index. FE refers to fixed effects estimation, RE refers to random effects estimation, se refers to the standard error, and r indicates that standard errors are robust.

Table 3.A.5: Contribution of the 2-digit sectors to the within-component of inequality

	(1)		(2)	(3)		(4)	
	Within-sectoral inequality		Weight (=sectoral wage share)	Weighted within-sectoral inequality		Subcategories per sector (average)	
sector	mean	rank	mean	mean	rank	3 digit	4 digit
15	0.0575	1	0.209	0.01477	1	4.8	12.9
16	0	22	0.0218	0	22	1	1
17	0.0182	11	0.0934	0.00094	11	2.8	5.6
18	0.0019	21	0.0739	0.00008	20	1.6	1.6
19	0.0165	15	0.0195	0.0003	18	1.9	2.5
20	0.0175	13	0.0355	0.00065	13	1.9	4.1
21	0.0115	19	0.0294	0.00037	17	1	2.7
22	0.0232	5	0.0451	0.0012	7	2.5	5.3
23	0.0136	17	0.0315	0.00057	15	1.5	1.5
24	0.0254	4	0.0743	0.00284	3	2.5	7
25	0.0138	16	0.0376	0.0006	14	1.9	2.5
26	0.0486	2	0.0615	0.00474	2	1.9	6.1
27	0.0178	12	0.057	0.00107	10	2.6	3
28	0.0188	10	0.0611	0.00114	9	1.9	6
29	0.0224	6	0.0605	0.00122	6	2.9	11.9
30	0	23	0.0148	0	23	1	1
31	0.0193	9	0.0468	0.00056	16	5.2	5.2
32	0.0202	8	0.039	0.00066	12	2.7	2.7
33	0.0128	18	0.0138	0.00022	19	2.4	3.9
34	0.0217	7	0.0544	0.00117	8	2.5	2.5
35	0.027	3	0.0333	0.00129	5	3.3	4.8
36	0.0166	14	0.0367	0.00176	4	1.9	4.3
37	0.0105	20	0.0026	0.00004	21	1.7	1.7

Notes. Columns (1) and (2) contain the unweighted and weighted within-components and the ranking of every 2-digit sector for each of these categories. Column (2) contains the weight and links the numbers in columns (1) and (2). Columns (4) display the average number of sectors covered by the data at the 3- and 4-digit level.

Table 3.A.6: Sectoral composition vs. sectoral coverage: FE and RE results

	(1)		(2)	
	Fixed effects	Standard error	Random effects	Standard error
Subsectors	0.0261***	(0.009)	0.0245***	(0.006)
Share_15	0.799	(0.970)	0.219	(0.692)
Share_17	-0.0140	(0.950)	-0.209	(0.633)
Share_18	-0.616	(0.951)	-0.986	(0.653)
Share_19	-0.654	(1.306)	-0.520	(1.032)
Share_20	2.989**	(1.416)	1.057	(0.865)
Share_21	0.370	(1.401)	-1.124	(0.843)
Share_22	-0.452	(0.967)	-0.257	(0.790)
Share_23	-1.629	(1.508)	-1.089	(1.061)
Share_24	-0.542	(0.976)	-0.725	(0.695)
Share_25	0.604	(1.037)	0.640	(0.911)
Share_26	1.135	(1.040)	0.735	(0.649)
Share_27	-1.132	(0.817)	-1.301**	(0.636)
Share_28	0.658	(0.866)	-0.00429	(0.736)
Share_29	-0.181	(0.900)	-0.476	(0.754)
Share_31	-0.489	(1.231)	-0.761	(0.904)
Share_32	-0.578	(0.902)	-0.788	(0.764)
Share_33	-0.108	(1.516)	1.158	(1.376)
Share_34	-0.446	(1.028)	-1.049*	(0.618)
Share_35	-0.567	(0.877)	-0.451	(0.657)
Share_36	-2.478*	(1.297)	-2.267*	(1.232)
Share_37	2.189	(2.656)	-0.343	(1.824)
Constant	26.44	(85.10)	57.99	(62.20)
Year FE	YES		YES	
Observations	429		429	
R^2	0.465			
# of countries	53		53	

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the share of within-sectoral inequality in percent. Numbers 15 to 37 refer to the 2-digit sector's wage share in total manufacturing wages and are also in percent. High variation samples 1 and to refer to subsamples of countries with above-average variation in sectoral coverage (1), and countries with one standard deviation above the average variation (2). Note that a Hausman test clearly rejects the random effects model at the <1% significance level.

Table 3.A.7: Relationship between Theil and income inequality: FE results, extended model

	(1) wiid	(2) swiid_n	(3) swiid_m	(4) lis	(5) silc	(6) wb
ln(Theil)	0.0168	0.00776	0.00521	-0.00260	0.0272	0.00956
	(0.0123)	(0.0128)	(0.0136)	(0.0323)	(0.0166)	(0.0177)
GDPpc	-9.46e-07	1.42e-06	9.69e-07	-4.10e-06	-2.30e-06	1.23e-06
	(3.02e-06)	(2.02e-06)	(2.50e-06)	(3.58e-06)	(2.86e-06)	(5.35e-06)
Population growth	-5.58e-05	0.00505	0.00727	0.0121	-0.0185	-0.00472
	(0.0130)	(0.0101)	(0.0110)	(0.0347)	(0.0131)	(0.0169)
Share urban	-0.00950**	0.00206	-0.00148	0.0205*	-0.00365	0.00770**
	(0.00408)	(0.00312)	(0.00374)	(0.0118)	(0.00512)	(0.00314)
Manuf. value added	-0.00683*	-0.00323*	-0.00430**	-0.00439	-0.00244	0.000585
	(0.00356)	(0.00188)	(0.00196)	(0.00773)	(0.00593)	(0.00334)
Trade openness	0.000746*	0.000148	2.53e-05	0.00101	0.000423	0.000440
	(0.000447)	(0.000310)	(0.000364)	(0.000802)	(0.000526)	(0.000494)
Price level of inv.	0.0217	0.0463	0.0589	-0.116	0.107	0.0839
	(0.0559)	(0.0328)	(0.0366)	(0.0919)	(0.0920)	(0.0619)
# imputed	-0.00244**	0.000900	0.00118	-0.00306	0.00157	-0.00101
	(0.00120)	(0.00123)	(0.00156)	(0.00320)	(0.00152)	(0.00174)
Tfp	0.264**	0.0612	0.0780	0.329	0.321**	0.210*
	(0.110)	(0.0501)	(0.0719)	(0.248)	(0.117)	(0.113)
1964.year		0.0353	0.0314			
		(0.0326)	(0.0285)			
1965.year		0.0124	0.00389			
		(0.0363)	(0.0293)			
1966.year		-0.0156	-0.0181			
		(0.0192)	(0.0260)			
1967.year		0.0113	-0.00363			
		(0.0350)	(0.0340)			
1968.year		-0.0148	-0.0298			
		(0.0152)	(0.0237)			
1969.year		-0.000286	-0.0269			
		(0.0284)	(0.0377)			
1970.year	0.490***	0.0133	0.00139			
	(0.0907)	(0.0269)	(0.0231)			
1971.year		-0.00819	0.000146			
		(0.0281)	(0.0308)			
1972.year	0.409***	-0.0242	-0.00968			
	(0.108)	(0.0240)	(0.0421)			
1973.year	0.626***	-0.0141	-0.00184			
	(0.0859)	(0.0239)	(0.0402)			
1974.year		-0.00593	-0.00541			
		(0.0244)	(0.0451)			
1975.year	-0.0259	-0.0137	-0.0276			
	(0.0297)	(0.0254)	(0.0439)			
1976.year	0.491***	-0.0143	-0.0148			
	(0.100)	(0.0242)	(0.0450)			
1977.year		-0.00480	-0.0123			
		(0.0281)	(0.0443)			
1978.year		-0.00750	-0.00795			
		(0.0393)	(0.0508)			
1979.year	0.343***	-0.00225	-0.0101			
	(0.0797)	(0.0412)	(0.0478)			
1980.year	0.375***	-0.0114	-0.0307			
	(0.0961)	(0.0476)	(0.0533)			
1981.year		-0.0310	-0.0452			
		(0.0442)	(0.0554)			
1982.year	0.262***	-0.0361	-0.0451			
	(0.0886)	(0.0417)	(0.0561)			
1983.year	0.0258	-0.0286	-0.0392			-0.0656
	(0.115)	(0.0379)	(0.0547)			(0.0744)
1984.year	0.319***	-0.0449	-0.0506	0.0381		-0.0814
	(0.0954)	(0.0374)	(0.0537)	(0.0686)		(0.0658)
1985.year	0.207*	-0.0447	-0.0491			-0.0870
	(0.120)	(0.0361)	(0.0501)			(0.0841)
1986.year	0.246***	-0.0494	-0.0458	0.00666		-0.104
	(0.0629)	(0.0337)	(0.0487)	(0.0343)		(0.105)
1987.year	0.292***	-0.0585	-0.0535			-0.0425
	(0.0851)	(0.0370)	(0.0513)			(0.0582)
1988.year	0.329***	-0.0583	-0.0465			-0.0873
	(0.0826)	(0.0391)	(0.0534)			(0.0811)

1989.year	0.291***	-0.0638	-0.0447	0.0680		-0.0436
	(0.0691)	(0.0415)	(0.0558)	(0.0436)		(0.0821)
1990.year	0.322***	-0.0715	-0.0460			-0.0986
	(0.0744)	(0.0479)	(0.0606)			(0.0800)
1991.year	0.213**	-0.0519	-0.0248	-0.0291		-0.0708
	(0.0867)	(0.0476)	(0.0610)	(0.0574)		(0.0830)
1992.year	0.284***	-0.0502	-0.0188	0.0772		-0.101
	(0.0748)	(0.0488)	(0.0621)	(0.0524)		(0.0878)
1993.year	0.261***	-0.0401	0.000750	0.0759		-0.0755
	(0.0786)	(0.0496)	(0.0631)	(0.0524)		(0.0885)
1994.year	0.330***	-0.0393	0.00168	0.0710		-0.120
	(0.0760)	(0.0508)	(0.0639)	(0.0623)		(0.0777)
1995.year	0.342***	-0.0312	0.00722	0.0525		-0.0759
	(0.0761)	(0.0526)	(0.0645)	(0.0505)		(0.0819)
1996.year	0.319***	-0.0290	0.00663	0.0497	-0.0181	-0.0803
	(0.0722)	(0.0538)	(0.0657)	(0.0521)	(0.0180)	(0.0801)
1997.year	0.341***	-0.0299	0.00524	0.0368	-0.0455**	-0.0911
	(0.0754)	(0.0554)	(0.0672)	(0.0507)	(0.0208)	(0.0797)
1998.year	0.350***	-0.0333	0.00134	0.101*	-0.0341	-0.0908
	(0.0732)	(0.0559)	(0.0682)	(0.0532)	(0.0276)	(0.0803)
1999.year	0.346***	-0.0355	-0.00411	0.0406	-0.0331	-0.0796
	(0.0758)	(0.0555)	(0.0682)	(0.0529)	(0.0236)	(0.0814)
2000.year	0.325***	-0.0303	-0.000648	0.0384	-0.0153	-0.0903
	(0.0748)	(0.0564)	(0.0701)	(0.0629)	(0.0349)	(0.0814)
2001.year	0.327***	-0.0315	0.000166	-0.0361	-0.0162	-0.0809
	(0.0712)	(0.0570)	(0.0709)	(0.104)	(0.0356)	(0.0826)
2002.year	0.348***	-0.0362	-0.00356	0.0416	-0.0269	-0.0830
	(0.0795)	(0.0581)	(0.0725)	(0.0555)	(0.0387)	(0.0848)
2003.year	0.324***	-0.0419	-0.00447	0.0243	-0.00712	-0.0845
	(0.0792)	(0.0594)	(0.0741)	(0.0775)	(0.0408)	(0.0829)
2004.year	0.338***	-0.0463	-0.00648	0.0499	-0.0360	-0.111
	(0.0846)	(0.0601)	(0.0760)	(0.0708)	(0.0433)	(0.0825)
2005.year	0.337***	-0.0494	-0.0101	0.0327	-0.0152	-0.126
	(0.0866)	(0.0606)	(0.0774)	(0.0760)	(0.0424)	(0.0832)
2006.year	0.312***	-0.0585	-0.0179	-0.194	-0.0207	-0.142
	(0.0935)	(0.0619)	(0.0790)	(0.163)	(0.0449)	(0.0875)
2007.year	0.315***	-0.0578	-0.0158	0.00534	-0.0522	-0.152*
	(0.0937)	(0.0647)	(0.0822)	(0.0899)	(0.0514)	(0.0899)
2008.year	0.329***	-0.0747	-0.0279	0.143	-0.0611	-0.172*
	(0.0921)	(0.0650)	(0.0819)	(0.114)	(0.0515)	(0.0920)
2009.year	0.364***	-0.0637	-0.0108	-0.112	-0.0307	-0.145
	(0.0839)	(0.0633)	(0.0789)	(0.132)	(0.0520)	(0.0877)
2010.year	0.311***	-0.0429	-0.00323	0.0502	-0.0353	-0.199**
	(0.0890)	(0.0660)	(0.0801)	(0.0782)	(0.0448)	(0.0935)
Constant	3.923***	3.479***	3.887***	-2.838***	3.461***	2.991***
	(0.324)	(0.132)	(0.181)	(0.964)	(0.417)	(0.279)
WIID dummmies	YES	-	-	-	-	
Observations	620	1,521	1,511	120	256	483
R^2	0.81	0.12	0.166	0.522	0.25	0.188
# of countries	66	82	82	35	28	73

Notes. Standard errors in parenthese; *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the logged Gini coefficient from the data source indicated in the top row. Swiid_n and swiid_m refer to net and market inequality, respectively. Silc denotes the EU SILC data and wb the WDI Gini coefficients. The dummies for the underlying WIID categories are included in column 1 but are not shown to save space (available upon request).

Table 3.A.8: Relationship between Theil and income inequality: RE results, reduced model

	(1) wiid	(2) swiid_n	(3) swiid_m	(4) lis	(5) silc	(6) wb
ln(Theil)	0.0282**	0.00900	0.00390	0.0378	0.0799**	0.00679
	(0.0124)	(0.0128)	(0.0133)	(0.0259)	(0.0325)	(0.0161)
GDPpc	-2.46e-06	3.24e-07	1.15e-06	-1.96e-06	-4.05e-06	-3.13e-06
	(2.50e-06)	(1.90e-06)	(2.35e-06)	(3.45e-06)	(3.59e-06)	(2.10e-06)
Population growth	0.0103	0.00623	0.00811	0.101***	0.0276	0.00265
	(0.0139)	(0.00962)	(0.0101)	(0.0319)	(0.0378)	(0.0132)
Share urban	-0.00347**	0.000770	-0.000387	-0.000549	-0.000786	0.00159
	(0.00155)	(0.00223)	(0.00250)	(0.00234)	(0.00248)	(0.00104)
Manuf. value added	-0.00698**	-0.00245	-0.00464***	-0.0102**	-0.0120***	0.000545
	(0.00317)	(0.00154)	(0.00158)	(0.00424)	(0.00459)	(0.00261)
Trade openness	0.000775**	0.000115	8.93e-05	-0.000461	-0.000414	6.48e-05
	(0.000315)	(0.000243)	(0.000268)	(0.000636)	(0.000588)	(0.000293)
Price level of inv.	0.00102	0.0447	0.0537	-0.444***	-0.194	0.0193
	(0.0431)	(0.0350)	(0.0364)	(0.137)	(0.189)	(0.0559)
# imputed	-0.00128	0.00106	0.00177	0.00447	0.00176	-0.00226
	(0.00119)	(0.00110)	(0.00152)	(0.00514)	(0.00518)	(0.00175)
Tfp	0.313***	0.0634	0.0800	0.671***	0.566***	0.226**
	(0.108)	(0.0514)	(0.0697)	(0.218)	(0.176)	(0.110)
# of ISIC	0.00626	0.00124	0.00351	0.0195*	0.00264	-0.00389
	(0.00472)	(0.00316)	(0.00297)	(0.0108)	(0.00996)	(0.00418)
ECA	-0.177***	-0.204***	-0.00948	-0.0234		-0.149**
	(0.0566)	(0.0586)	(0.0551)	(0.116)		(0.0599)
LAC	0.288***	0.235***	0.0947	0.154		0.257***
	(0.0672)	(0.0669)	(0.0593)	(0.119)		(0.0662)
MENA	-0.125	-0.00518	-0.0365		-0.143	0.0274
	(0.0951)	(0.0913)	(0.0652)		(0.0885)	(0.0815)
NA	-0.0649	-0.0950	0.0246	0.127		0.0720
	(0.0931)	(0.0870)	(0.0802)	(0.139)		(0.0869)
SA	-0.102	0.173	0.0246	-0.0358		-0.0549
	(0.154)	(0.138)	(0.104)	(0.164)		(0.0776)
SSH	0.116	0.239***	0.124	0.0139		0.217**
	(0.137)	(0.0819)	(0.0836)	(0.140)		(0.0903)
1964.year		0.0380	0.0304			
		(0.0352)	(0.0277)			
1965.year		0.0111	0.00700			
		(0.0370)	(0.0273)			
1966.year		-0.0166	-0.0163			
		(0.0205)	(0.0234)			
1967.year		0.0132	-0.00289			
		(0.0374)	(0.0318)			
1968.year		-0.0124	-0.0308			
		(0.0172)	(0.0214)			
1969.year		0.00421	-0.0287			
		(0.0312)	(0.0356)			
1970.year	0.474***	0.0174	-0.00349			
	(0.0746)	(0.0283)	(0.0210)			
1971.year		-0.00249	-0.00396			
		(0.0292)	(0.0281)			
1972.year	0.505***	-0.0200	-0.0141			
	(0.0602)	(0.0238)	(0.0401)			
1973.year	0.603***	-0.00738	-0.00841			
	(0.0734)	(0.0237)	(0.0370)			
1974.year		0.00346	-0.0134			
		(0.0231)	(0.0404)			
1975.year	-0.0496**	-0.00203	-0.0358			
	(0.0230)	(0.0214)	(0.0370)			
1976.year	0.509***	-0.00364	-0.0228			
	(0.0587)	(0.0201)	(0.0386)			
1977.year		0.00790	-0.0207			
		(0.0219)	(0.0365)			
1978.year		0.00513	-0.0170			
		(0.0325)	(0.0422)			
1979.year	0.224***	0.0122	-0.0200			
	(0.0802)	(0.0326)	(0.0377)			
1980.year	0.361***	0.00488	-0.0414			
	(0.118)	(0.0392)	(0.0416)			
1981.year		-0.0150	-0.0560			
		(0.0344)	(0.0435)			
1982.year	0.224***	-0.0188	-0.0569			
	(0.0855)	(0.0302)	(0.0428)			

1983.year	0.0408	-0.0107	-0.0512			-0.0102
	(0.118)	(0.0271)	(0.0410)			(0.101)
1984.year	0.298***	-0.0268	-0.0627	-0.0767		-0.0708
	(0.0993)	(0.0269)	(0.0404)	(0.0900)		(0.0834)
1985.year	0.174	-0.0265	-0.0611*			-0.0442
	(0.118)	(0.0260)	(0.0367)			(0.104)
1986.year	0.206***	-0.0306	-0.0581*	-0.00174		-0.0779
	(0.0580)	(0.0241)	(0.0353)	(0.0173)		(0.125)
1987.year	0.280***	-0.0392	-0.0667*			-0.00218
	(0.0712)	(0.0270)	(0.0373)			(0.0805)
1988.year	0.291***	-0.0377	-0.0608			-0.0135
	(0.0746)	(0.0290)	(0.0382)			(0.106)
1989.year	0.258***	-0.0417	-0.0590	0.0501		0.00678
	(0.0634)	(0.0309)	(0.0397)	(0.0651)		(0.107)
1990.year	0.284***	-0.0484	-0.0622			-0.0339
	(0.0641)	(0.0365)	(0.0433)			(0.104)
1991.year	0.170**	-0.0280	-0.0418	0.118*		0.00235
	(0.0842)	(0.0363)	(0.0429)	(0.0704)		(0.108)
1992.year	0.242***	-0.0254	-0.0361	0.153***		-0.0289
	(0.0671)	(0.0386)	(0.0447)	(0.0578)		(0.113)
1993.year	0.211***	-0.0138	-0.0177	0.287***		0.00147
	(0.0701)	(0.0378)	(0.0443)	(0.0560)		(0.116)
1994.year	0.285***	-0.0121	-0.0172	0.209***		-0.0309
	(0.0672)	(0.0383)	(0.0445)	(0.0556)		(0.102)
1995.year	0.298***	-0.00373	-0.0119	0.225***		0.0148
	(0.0664)	(0.0396)	(0.0450)	(0.0457)		(0.108)
1996.year	0.274***	-0.000551	-0.0123	0.181***	-0.00717	0.00806
	(0.0616)	(0.0403)	(0.0456)	(0.0482)	(0.0274)	(0.103)
1997.year	0.294***	-0.000578	-0.0140	0.178***	-0.0700***	0.0120
	(0.0671)	(0.0414)	(0.0466)	(0.0523)	(0.0259)	(0.102)
1998.year	0.302***	-0.00277	-0.0184	0.294***	-0.0185	0.00709
	(0.0615)	(0.0414)	(0.0470)	(0.0619)	(0.0347)	(0.104)
1999.year	0.299***	-0.00422	-0.0252	0.120	-0.0425	0.0205
	(0.0657)	(0.0405)	(0.0464)	(0.0851)	(0.0269)	(0.104)
2000.year	0.273***	0.00229	-0.0219	0.172***	-0.0431	0.0154
	(0.0660)	(0.0409)	(0.0475)	(0.0547)	(0.0294)	(0.100)
2001.year	0.275***	0.00207	-0.0216	-0.0524	-0.0602**	0.0196
	(0.0617)	(0.0411)	(0.0478)	(0.145)	(0.0279)	(0.100)
2002.year	0.301***	-0.00154	-0.0247	0.375*	-0.0667	0.0325
	(0.0709)	(0.0420)	(0.0492)	(0.199)	(0.0456)	(0.102)
2003.year	0.274***	-0.00589	-0.0255	-0.00139	0.00136	0.0305
	(0.0679)	(0.0431)	(0.0503)	(0.157)	(0.0414)	(0.0983)
2004.year	0.291***	-0.00909	-0.0278	0.264***	-0.00221	0.00957
	(0.0725)	(0.0437)	(0.0517)	(0.0559)	(0.0546)	(0.0994)
2005.year	0.285***	-0.0110	-0.0321	0.158**	-0.0161	0.00142
	(0.0736)	(0.0435)	(0.0522)	(0.0772)	(0.0463)	(0.101)
2006.year	0.259***	-0.0189	-0.0404	0.140	-0.00836	-0.00822
	(0.0794)	(0.0443)	(0.0533)	(0.131)	(0.0511)	(0.101)
2007.year	0.261***	-0.0164	-0.0381	0.256***	-0.0130	-0.00417
	(0.0771)	(0.0470)	(0.0567)	(0.0798)	(0.0706)	(0.101)
2008.year	0.284***	-0.0308	-0.0472	0.472***	0.0110	-0.0130
	(0.0775)	(0.0471)	(0.0569)	(0.130)	(0.0857)	(0.101)
2009.year	0.323***	-0.0192	-0.0261	0.265***	-0.00436	-0.00359
	(0.0734)	(0.0458)	(0.0555)	(0.0953)	(0.0696)	(0.100)
2010.year	0.276***	0.00240	-0.0202	0.322***	-0.00353	-0.0498
	(0.0794)	(0.0510)	(0.0589)	(0.0855)	(0.0688)	(0.103)
Constant	3.562***	3.503***	3.738***	-1.635***	3.814***	-1.779***
	(0.180)	(0.129)	(0.136)	(0.372)	(0.110)	(0.520)
WIID dummmies	YES	-	-	-	-	-
Observations	620	1,521	1,511	120	1,741	121
R^2						0.481
# of countries	66	82	82	35	100	36

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the logged Gini coefficient from the data source indicated in the top row. Swiid_n and swiid_m refer to net and market inequality, respectively. Silc denotes the EU SILC data and wb the WDI Gini coefficients. The dummies for the underlying WIID categories are included in column 1 but are not shown to save space (available upon request).

Table 3.A.9: Relationship between Theil and income inequality: FE results, reduced model

	(1) wiid	(2) swiid_n	(3) swiid_m	(4) lis	(5) silc	(6) wb
ln(Theil)	0.00909	0.0115	0.00788	0.00468	0.0315	0.00354
	(0.0112)	(0.0117)	(0.0122)	(0.0312)	(0.0187)	(0.0146)
Population growth	-0.00217	0.0134**	0.0150*	-0.0252	-0.0287**	0.00701
	(0.0205)	(0.00652)	(0.00790)	(0.0247)	(0.0133)	(0.0163)
Share urban	-0.0134***	0.00270	-0.000306	0.00753	-0.00758*	0.00802*
	(0.00403)	(0.00294)	(0.00337)	(0.00737)	(0.00427)	(0.00424)
Manuf. value added	-0.00522	-0.00307*	-0.00416**	0.00177	0.00189	-0.00120
	(0.00377)	(0.00176)	(0.00199)	(0.00574)	(0.00598)	(0.00335)
1964.year		0.0183	0.0224			
		(0.0206)	(0.0222)			
1965.year		-0.00567	-0.0159			
		(0.0233)	(0.0242)			
1966.year		-0.0261*	-0.0318			
		(0.0155)	(0.0242)			
1967.year		-0.0188	-0.0224			
		(0.0263)	(0.0311)			
1968.year		0.00719	-0.00941			
		(0.0322)	(0.0297)			
1969.year		-0.0144	-0.0243			
		(0.0262)	(0.0242)			
1970.year	0.484***	-0.00146	-0.0139			
	(0.0693)	(0.0291)	(0.0248)			
1971.year		-0.0265	-0.0198			
		(0.0277)	(0.0275)			
1972.year	0.357***	-0.0328	-0.0253			
	(0.102)	(0.0311)	(0.0364)			
1973.year	0.643***	-0.0319	-0.0256			
	(0.0689)	(0.0272)	(0.0350)			
1974.year		-0.0212	-0.0255			
		(0.0286)	(0.0355)			
1975.year	0.0349**	-0.0327	-0.0396			
	(0.0166)	(0.0289)	(0.0326)			
1976.year	0.510***	-0.0271	-0.0325			
	(0.0981)	(0.0264)	(0.0353)			
1977.year		-0.0115	-0.0249			
		(0.0295)	(0.0342)			
1978.year		-0.0141	-0.0200			
		(0.0409)	(0.0431)			
1979.year	0.336***	-0.0110	-0.0234			
	(0.0707)	(0.0423)	(0.0424)			
1980.year	0.418***	-0.0182	-0.0410			
	(0.0809)	(0.0466)	(0.0460)			
1981.year		-0.0387	-0.0646			
		(0.0403)	(0.0463)			
1982.year	0.342***	-0.0499	-0.0716			
	(0.0768)	(0.0419)	(0.0500)			
1983.year	0.313***	-0.0341	-0.0541			-0.0979
	(0.0510)	(0.0396)	(0.0467)			(0.125)
1984.year	0.410***	-0.0529	-0.0636	0.0670		-0.0655
	(0.0828)	(0.0385)	(0.0456)	(0.0452)		(0.102)
1985.year	0.300***	-0.0482	-0.0547			-0.101
	(0.109)	(0.0385)	(0.0441)			(0.133)
1986.year	0.315***	-0.0494	-0.0571	0.0416*		-0.0755
	(0.0520)	(0.0402)	(0.0457)	(0.0242)		(0.143)
1987.year	0.376***	-0.0536	-0.0597			-0.0241
	(0.0795)	(0.0434)	(0.0479)			(0.101)
1988.year	0.421***	-0.0516	-0.0494			-0.0488
	(0.0748)	(0.0448)	(0.0492)			(0.127)
1989.year	0.375***	-0.0587	-0.0425	0.0924**		-0.00625
	(0.0672)	(0.0468)	(0.0513)	(0.0431)		(0.131)
1990.year	0.415***	-0.0717	-0.0475			-0.0595
	(0.0701)	(0.0487)	(0.0536)			(0.122)

1991.year	0.324***	-0.0512	-0.0244	-0.0149		-0.0340
	(0.0744)	(0.0490)	(0.0545)	(0.0519)		(0.129)
1992.year	0.393***	-0.0456	-0.0136	0.110*		-0.0616
	(0.0644)	(0.0509)	(0.0562)	(0.0573)		(0.132)
1993.year	0.370***	-0.0368	0.000175	0.122**		-0.0402
	(0.0660)	(0.0528)	(0.0580)	(0.0523)		(0.131)
1994.year	0.447***	-0.0358	-0.00162	0.117*		-0.0820
	(0.0631)	(0.0532)	(0.0582)	(0.0651)		(0.125)
1995.year	0.449***	-0.0284	0.00608	0.0831*		0.00888
	(0.0658)	(0.0547)	(0.0594)	(0.0413)		(0.126)
1996.year	0.436***	-0.0303	-0.00113	0.0804	-0.0120	-0.0542
	(0.0597)	(0.0557)	(0.0604)	(0.0630)	(0.0170)	(0.124)
1997.year	0.463***	-0.0322	-0.00300	0.0542	-0.0433**	-0.0478
	(0.0649)	(0.0573)	(0.0614)	(0.0501)	(0.0193)	(0.126)
1998.year	0.480***	-0.0362	-0.00756	0.138**	-0.0272	-0.0653
	(0.0609)	(0.0582)	(0.0624)	(0.0519)	(0.0247)	(0.126)
1999.year	0.474***	-0.0408	-0.0169	0.0401	-0.0267	-0.0582
	(0.0652)	(0.0590)	(0.0631)	(0.0563)	(0.0164)	(0.128)
2000.year	0.465***	-0.0381	-0.0168	0.0964*	-0.0169	-0.0539
	(0.0658)	(0.0590)	(0.0634)	(0.0556)	(0.0270)	(0.126)
2001.year	0.466***	-0.0460	-0.0230	0.100	-0.0184	-0.0182
	(0.0608)	(0.0599)	(0.0650)	(0.0942)	(0.0280)	(0.126)
2002.year	0.485***	-0.0434	-0.0200	0.0599	-0.0226	-0.0530
	(0.0680)	(0.0611)	(0.0663)	(0.0673)	(0.0309)	(0.131)
2003.year	0.472***	-0.0454	-0.0170	0.0856	0.0277	-0.0432
	(0.0646)	(0.0617)	(0.0665)	(0.0656)	(0.0313)	(0.129)
2004.year	0.500***	-0.0459	-0.0163	0.0906	0.0115	-0.0531
	(0.0711)	(0.0621)	(0.0675)	(0.0632)	(0.0325)	(0.129)
2005.year	0.501***	-0.0446	-0.0163	0.0789	0.0337	-0.0591
	(0.0697)	(0.0624)	(0.0683)	(0.0687)	(0.0334)	(0.130)
2006.year	0.488***	-0.0454	-0.0169	-0.148	0.0377	-0.0553
	(0.0747)	(0.0622)	(0.0682)	(0.154)	(0.0354)	(0.130)
2007.year	0.495***	-0.0333	-0.00160	0.0606	0.0225	-0.0535
	(0.0724)	(0.0634)	(0.0686)	(0.0740)	(0.0379)	(0.129)
2008.year	0.505***	-0.0520	-0.0126	0.165	0.0195	-0.0738
	(0.0705)	(0.0656)	(0.0696)	(0.122)	(0.0399)	(0.130)
2009.year	0.519***	-0.0502	-0.00256	-0.0644	0.0249	-0.0750
	(0.0705)	(0.0675)	(0.0709)	(0.125)	(0.0483)	(0.130)
2010.year	0.471***	-0.0125	0.0221	0.0720	0.0166	-0.102
	(0.0739)	(0.0658)	(0.0688)	(0.0811)	(0.0423)	(0.133)
Constant	4.421***	3.566***	3.927***	-1.779***	3.992***	3.489***
	(0.301)	(0.0943)	(0.126)	(0.520)	(0.200)	(0.201)
Observations	633	1,765	1,741	121	256	538
WIID dummmies	YES	-	-	-	-	-
R^2	0.799	0.057	0.089	0.481		
# of countries	71	100	100	36	28	87

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the logged Gini coefficient from the data source indicated in the top row. Swiid_n and swiid_m refer to net and market inequality, respectively. Silc denotes the EU SILC data and wb the WDI Gini coefficients. The dummies for the underlying WIID categories are included in column 2 but are not shown to save space (available upon request).

Table 3.A.10: Determinants of the difference between wage and income inequality, FE results

	(1) wiid	(2) swiid_n	(3) swiid_m	(4) lis	(5) silc	(6) wb
GDPpc	9.51e-06	1.78e-05	1.77e-05	2.72e-05	2.05e-05	4.40e-05
	(1.03e-05)	(1.25e-05)	(1.24e-05)	(1.97e-05)	(1.40e-05)	(3.07e-05)
Population growth	0.192**	0.0493	0.0488	-0.0548	-0.0392	0.0992
	(0.0740)	(0.0503)	(0.0504)	(0.151)	(0.0750)	(0.0724)
Share urban	0.0304*	0.00710	0.00224	-0.0243	-0.0251	0.0200
	(0.0169)	(0.0127)	(0.0119)	(0.0764)	(0.0274)	(0.0153)
Manuf. value added	0.00266	0.0173*	0.0193**	0.0307	0.0247	0.0146
	(0.0147)	(0.00964)	(0.00910)	(0.0405)	(0.0260)	(0.0152)
Trade openness	-0.00352	-0.00471**	-0.00481**	-0.00318	-0.00278	-0.00389
	(0.00289)	(0.00234)	(0.00241)	(0.00301)	(0.00255)	(0.00304)
Price level of inv.	-0.470*	0.485	0.522	0.324	-0.473	-0.0473
	(0.271)	(0.327)	(0.330)	(0.580)	(0.680)	(0.301)
# imputed	-0.0408**	-0.0659***	-0.0659***	0.000215	-0.0355**	-0.0649***
	(0.0172)	(0.0181)	(0.0185)	(0.0136)	(0.0138)	(0.0200)
Tfp	0.895*	-0.513***	-0.460*	0.257	0.349	-0.198
	(0.459)	(0.192)	(0.249)	(1.154)	(0.659)	(0.416)
1964.year		-0.139	-0.120			
		(0.160)	(0.214)			
1965.year		-0.0412	-0.00376			
		(0.213)	(0.257)			
1966.year		-0.0591	-0.0361			
		(0.185)	(0.239)			
1967.year		-0.0481	-0.0608			
		(0.216)	(0.264)			
1968.year		-0.0740	-0.0737			
		(0.184)	(0.247)			
1969.year		-0.328	-0.364			
		(0.228)	(0.305)			
1970.year	1.610***	-0.00482	-0.0167			
	(0.437)	(0.246)	(0.297)			
1971.year		-0.227	-0.221			
		(0.234)	(0.301)			
1972.year	1.540***	-0.223	-0.214			
	(0.448)	(0.209)	(0.272)			
1973.year	1.751***	-0.0708	-0.0644			
	(0.428)	(0.218)	(0.287)			
1974.year		-0.165	-0.174			
		(0.240)	(0.306)			
1975.year	-0.147	-0.172	-0.190			
	(0.0888)	(0.262)	(0.322)			
1976.year	1.573***	-0.269	-0.277			
	(0.366)	(0.255)	(0.315)			
1977.year		-0.199	-0.212			
		(0.254)	(0.300)			
1978.year		0.0223	0.0170			
		(0.280)	(0.337)			
1979.year	1.671***	-0.0736	-0.0855			
	(0.406)	(0.290)	(0.339)			
1980.year	1.536***	-0.132	-0.152			
	(0.327)	(0.279)	(0.329)			
1981.year		-0.0300	-0.0465			
		(0.269)	(0.327)			
1982.year	1.466***	-0.189	-0.195			
	(0.317)	(0.263)	(0.317)			
1983.year	1.806***	-0.145	-0.155			-1.909***
	(0.568)	(0.264)	(0.305)			(0.652)
1984.year	1.602***	-0.164	-0.172	0.137		-0.944*
	(0.305)	(0.257)	(0.298)	(0.545)		(0.518)
1985.year	1.523***	-0.140	-0.145			-1.770**
	(0.285)	(0.248)	(0.295)			(0.686)
1986.year	1.325***	-0.287	-0.284	0.189		-1.670**

	(0.300)	(0.253)	(0.299)	(0.228)		(0.746)
1987.year	1.316***	-0.473*	-0.471			-1.641**
	(0.341)	(0.267)	(0.308)			(0.628)
1988.year	1.060***	-0.459	-0.450			-2.160***
	(0.363)	(0.289)	(0.329)			(0.716)
1989.year	1.125***	-0.587*	-0.569	-0.615		-1.860***
	(0.346)	(0.303)	(0.343)	(0.446)		(0.698)
1990.year	1.330***	-0.611*	-0.588			-1.964**
	(0.321)	(0.331)	(0.359)			(0.771)
1991.year	0.951**	-0.614*	-0.566	-0.576		-2.062***
	(0.372)	(0.336)	(0.363)	(0.515)		(0.677)
1992.year	1.055***	-0.548	-0.494	-0.158		-1.921***
	(0.354)	(0.338)	(0.367)	(0.581)		(0.674)
1993.year	0.928***	-0.567*	-0.515	-0.561		-1.734**
	(0.317)	(0.340)	(0.368)	(0.497)		(0.730)
1994.year	1.078***	-0.571	-0.507	-0.309		-2.003***
	(0.348)	(0.347)	(0.374)	(0.557)		(0.691)
1995.year	1.141***	-0.560	-0.524	-0.496		-1.853***
	(0.327)	(0.353)	(0.382)	(0.485)		(0.661)
1996.year	1.010***	-0.633*	-0.599	-0.491	0.0249	-1.990***
	(0.364)	(0.361)	(0.387)	(0.521)	(0.0708)	(0.676)
1997.year	0.942**	-0.678*	-0.645*	-0.860*	-0.112	-1.928***
	(0.359)	(0.357)	(0.383)	(0.471)	(0.0868)	(0.697)
1998.year	0.907**	-0.740**	-0.706*	-0.635	-0.0915	-1.876***
	(0.346)	(0.354)	(0.378)	(0.496)	(0.0850)	(0.689)
1999.year	0.821**	-0.772**	-0.741*	-0.803	-0.0984	-2.118***
	(0.330)	(0.355)	(0.379)	(0.552)	(0.111)	(0.719)
2000.year	0.815**	-0.697*	-0.668*	-0.928*	-0.111	-2.044***
	(0.353)	(0.359)	(0.383)	(0.537)	(0.125)	(0.690)
2001.year	0.847**	-0.732**	-0.695*	-2.763***	-0.103	-2.215***
	(0.364)	(0.359)	(0.384)	(0.761)	(0.139)	(0.723)
2002.year	0.760**	-0.784**	-0.746*	-0.748	-0.402*	-2.172***
	(0.341)	(0.376)	(0.399)	(0.600)	(0.226)	(0.713)
2003.year	0.851**	-0.830**	-0.794*	-1.356*	0.0273	-2.097***
	(0.354)	(0.387)	(0.408)	(0.701)	(0.196)	(0.719)
2004.year	0.880**	-0.823**	-0.787*	-0.834	0.0672	-2.124***
	(0.385)	(0.404)	(0.422)	(0.711)	(0.219)	(0.727)
2005.year	0.854**	-0.809*	-0.772*	-1.330	0.0627	-2.134***
	(0.402)	(0.412)	(0.430)	(0.790)	(0.215)	(0.726)
2006.year	0.776*	-0.797*	-0.762*	-1.313	0.00280	-2.142***
	(0.415)	(0.429)	(0.447)	(0.818)	(0.236)	(0.726)
2007.year	0.912**	-0.883**	-0.850*	-0.939	0.108	-2.171***
	(0.438)	(0.434)	(0.449)	(0.816)	(0.287)	(0.727)
2008.year	0.830*	-0.857*	-0.810*	-0.942	0.0213	-2.116***
	(0.459)	(0.460)	(0.475)	(0.808)	(0.328)	(0.731)
2009.year	0.732	-0.910**	-0.850*	-0.921	-0.0215	-2.138***
	(0.484)	(0.457)	(0.470)	(0.867)	(0.246)	(0.733)
2010.year	0.902*	-0.691	-0.650	-0.765	0.113	-2.081***
	(0.460)	(0.464)	(0.475)	(0.828)	(0.261)	(0.755)
Constant	4.421***	8.101***	8.490***	4.682	9.531***	8.313***
	(1.173)	(0.737)	(0.718)	(5.502)	(2.312)	(0.982)
Observations	618	1,521	1,511	120	256	483
WIID dummmies	YES	-	-	-	-	-
R^2	0.383	0.327	0.323	0.560	0.228	0.358
# of countries	66	82	82	35	28	73

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the logged percentage difference between the (normalized) Theil index and the Gini coefficient from the data source indicated in the top row. Swiid_n and swiid_m refer to net and market inequality, respectively. Silc denotes the EU SILC data and wb the WDI Gini coefficients. The dummies for the underlying WIID categories are included in column 1 but are not shown to save space (available upon request).

Table 3.A.11: Determinants of the difference between wage and income inequality, RE results

	(1) wiid	(2) swiid_n	(3) swiid_m	(4) lis	(5) silc	(6) wb
GDP per capita	8.77e-06	1.65e-05*	2.02e-05**	2.22e-05**	-9.20e-07	2.69e-05**
	(7.29e-06)	(9.68e-06)	(1.01e-05)	(1.12e-05)	(1.13e-05)	(1.28e-05)
Population growth	0.146**	0.0613	0.0538	0.0371	0.0178	0.0734
	(0.0711)	(0.0437)	(0.0443)	(0.124)	(0.0986)	(0.0646)
Share urban	0.00199	0.00731	0.00791	0.00364	0.000452	0.00405
	(0.00436)	(0.00499)	(0.00514)	(0.0108)	(0.00825)	(0.00390)
Manufacturing value added	0.00633	0.0184**	0.0171**	0.0344*	0.00522	0.0168
	(0.0118)	(0.00818)	(0.00794)	(0.0191)	(0.0170)	(0.0130)
Trade openness	-0.000161	-0.00215*	-0.00234*	-0.00381*	-0.00276*	-0.00266*
	(0.00119)	(0.00123)	(0.00128)	(0.00201)	(0.00167)	(0.00143)
Price level of investment	0.0222	0.507*	0.546*	0.444	1.581***	0.0145
	(0.275)	(0.306)	(0.301)	(0.466)	(0.499)	(0.301)
# imputed	-0.0296**	-0.0603***	-0.0601***	0.00903	-0.0216	-0.0537**
	(0.0142)	(0.0204)	(0.0217)	(0.0106)	(0.0136)	(0.0242)
Total factor productivity	0.284	-0.496**	-0.464*	0.718	0.467	-0.0676
	(0.491)	(0.200)	(0.260)	(0.890)	(0.781)	(0.349)
# of ISIC	0.0728**	0.0297	0.0350	0.0619	0.00503	0.0465
	(0.0308)	(0.0317)	(0.0347)	(0.0379)	(0.0294)	(0.0354)
ECA	0.579**	0.257	0.397*	1.263***		0.273
	(0.258)	(0.229)	(0.223)	(0.372)		(0.262)
LAC	-0.0637	-0.188	-0.345*	0.492		-0.0266
	(0.264)	(0.210)	(0.201)	(0.607)		(0.269)
MENA	-0.201	-0.122	-0.159		-0.117	-0.179
	(0.212)	(0.271)	(0.271)		(0.329)	(0.425)
NA	0.386	0.00148	0.00384	0.931**		-0.0377
	(0.393)	(0.405)	(0.423)	(0.436)		(0.412)
SA	-0.245	0.101	0.0255	0.975		-0.113
	(0.315)	(0.330)	(0.367)	(0.650)		(0.388)
SSH	0.396	0.0207	-0.0352	1.536***		0.0944
	(0.472)	(0.271)	(0.271)	(0.480)		(0.365)
1964.year		-0.137	-0.124			
		(0.161)	(0.219)			
1965.year		-0.0312	0.0271			
		(0.214)	(0.250)			
1966.year		-0.0537	-0.0130			
		(0.186)	(0.231)			
1967.year		-0.0389	-0.0437			
		(0.218)	(0.258)			
1968.year		-0.0781	-0.0714			
		(0.184)	(0.238)			
1969.year		-0.315	-0.351			
		(0.229)	(0.306)			
1970.year	1.539***	-0.0234	-0.0382			
	(0.387)	(0.247)	(0.291)			
1971.year		-0.226	-0.229			
		(0.234)	(0.295)			
1972.year	0.886***	-0.238	-0.231			
	(0.224)	(0.205)	(0.264)			
1973.year	1.710***	-0.103	-0.108			
	(0.388)	(0.211)	(0.278)			
1974.year		-0.212	-0.241			
		(0.249)	(0.307)			
1975.year	-0.0739	-0.215	-0.260			
	(0.101)	(0.262)	(0.317)			
1976.year	1.261***	-0.312	-0.345			
	(0.379)	(0.253)	(0.310)			
1977.year		-0.239	-0.286			
		(0.245)	(0.286)			
1978.year		-0.00569	-0.0466			
		(0.266)	(0.323)			
1979.year	1.185**	-0.115	-0.168			
	(0.489)	(0.281)	(0.331)			
1980.year	1.125***	-0.190	-0.256			
	(0.270)	(0.261)	(0.311)			
1981.year		-0.0847	-0.145			
		(0.248)	(0.308)			
1982.year	1.339***	-0.226	-0.284			
	(0.358)	(0.236)	(0.293)			

1983.year	0.839	-0.192	-0.255			-1.914***
	(0.766)	(0.219)	(0.269)			(0.625)
1984.year	1.585***	-0.204	-0.266	0.0507		-0.978*
	(0.350)	(0.216)	(0.265)	(0.374)		(0.527)
1985.year	1.488***	-0.179	-0.237			-1.753***
	(0.320)	(0.212)	(0.266)			(0.638)
1986.year	1.226***	-0.311	-0.368	0.156*		-1.598**
	(0.336)	(0.214)	(0.267)	(0.0919)		(0.734)
1987.year	1.134***	-0.507**	-0.567**			-1.682***
	(0.337)	(0.229)	(0.274)			(0.616)
1988.year	0.967***	-0.502**	-0.561**			-2.105***
	(0.352)	(0.239)	(0.286)			(0.645)
1989.year	1.065***	-0.627**	-0.680**	-0.705**		-1.795***
	(0.341)	(0.255)	(0.302)	(0.277)		(0.669)
1990.year	1.254***	-0.663**	-0.717**			-1.935***
	(0.266)	(0.285)	(0.319)			(0.730)
1991.year	0.917**	-0.665**	-0.695**	-0.707**		-2.018***
	(0.397)	(0.288)	(0.321)	(0.304)		(0.631)
1992.year	1.083***	-0.590**	-0.622*	-0.306		-1.842***
	(0.360)	(0.290)	(0.326)	(0.342)		(0.625)
1993.year	0.980***	-0.619**	-0.658**	-0.713***		-1.692**
	(0.336)	(0.286)	(0.325)	(0.228)		(0.693)
1994.year	1.079***	-0.632**	-0.661**	-0.468*		-1.921***
	(0.364)	(0.289)	(0.325)	(0.241)		(0.636)
1995.year	1.108***	-0.624**	-0.681**	-0.641***		-1.816***
	(0.353)	(0.287)	(0.327)	(0.247)		(0.607)
1996.year	0.964**	-0.693**	-0.755**	-0.635***	-0.0446	-1.964***
	(0.395)	(0.289)	(0.329)	(0.177)	(0.0632)	(0.613)
1997.year	0.932**	-0.732**	-0.799**	-1.020***	0.0403	-1.804***
	(0.382)	(0.287)	(0.325)	(0.254)	(0.0869)	(0.633)
1998.year	0.910**	-0.794***	-0.866***	-0.783***	-0.0553	-1.767***
	(0.363)	(0.287)	(0.324)	(0.231)	(0.0839)	(0.626)
1999.year	0.830**	-0.833***	-0.912***	-1.006***	0.0380	-2.009***
	(0.351)	(0.284)	(0.325)	(0.299)	(0.128)	(0.651)
2000.year	0.852**	-0.770***	-0.852***	-1.075***	0.113	-1.938***
	(0.372)	(0.284)	(0.326)	(0.266)	(0.125)	(0.631)
2001.year	0.874**	-0.800***	-0.880***	-2.694***	0.172	-2.125***
	(0.374)	(0.285)	(0.330)	(0.370)	(0.137)	(0.660)
2002.year	0.766**	-0.844***	-0.925***	-0.823***	-0.356	-2.051***
	(0.339)	(0.290)	(0.336)	(0.255)	(0.227)	(0.631)
2003.year	0.838**	-0.886***	-0.973***	-1.412***	-0.210	-2.000***
	(0.352)	(0.295)	(0.337)	(0.363)	(0.162)	(0.638)
2004.year	0.838**	-0.886***	-0.976***	-1.024***	-0.294*	-2.016***
	(0.388)	(0.300)	(0.340)	(0.381)	(0.158)	(0.645)
2005.year	0.842**	-0.883***	-0.976***	-1.580***	-0.244	-2.004***
	(0.395)	(0.299)	(0.339)	(0.513)	(0.174)	(0.637)
2006.year	0.731*	-0.889***	-0.986***	-1.581***	-0.364*	-2.008***
	(0.400)	(0.312)	(0.352)	(0.512)	(0.192)	(0.644)
2007.year	0.815*	-0.978***	-1.081***	-1.189**	-0.440**	-2.025***
	(0.424)	(0.308)	(0.345)	(0.484)	(0.218)	(0.640)
2008.year	0.794*	-0.920***	-1.009***	-1.131**	-0.748***	-1.926***
	(0.439)	(0.330)	(0.364)	(0.527)	(0.257)	(0.634)
2009.year	0.843**	-0.904***	-0.978***	-1.081**	-0.584***	-1.884***
	(0.411)	(0.319)	(0.350)	(0.489)	(0.177)	(0.642)
2010.year	1.044**	-0.694**	-0.793**	-0.833	-0.136	-1.832***
	(0.443)	(0.334)	(0.364)	(0.605)	(0.225)	(0.643)
Constant	4.319***	7.279***	7.382***	0.0639	6.742***	8.160***
	(0.709)	(0.938)	(0.996)	(1.502)	(0.836)	(1.102)
Observations	618	1,521	1,511	120	256	483
WIID dummies	YES	-	-	-	-	-
R^2	0.383	0.327	0.323	0.560	0.228	0.358
# of countries	66	82	82	35	28	73

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the logged percentage difference between the (normalized) Theil index and the Gini coefficient from the data source indicated in the top row. Swiid_n and swiid_m refer to net and market inequality, respectively. Silc denotes the EU SILC data and wb the WDI Gini coefficients. The dummies for the underlying WIID categories are included in column 1 but are not shown to save space (available upon request).

Table 3.A.12: Relationship between Theil and Income Inequality: FE results, using the UTIP index

	(1) wiid	(2) swiid_n	(3) swiid_m	(4) lis	(5) silc	(6) wb
ln(Theil_utip)	0.0903***	0.0116	0.0162	0.0731**	0.0335	0.0218
	(0.0195)	(0.0178)	(0.0194)	(0.0273)	(0.0254)	(0.0171)
Pop. growth	-0.0179	0.0163***	0.0155*	-0.0158	-0.0266	0.00745
	(0.0206)	(0.00581)	(0.00801)	(0.0255)	(0.0181)	(0.0173)
Share urban	-0.0135***	0.00168	-0.000666	0.00304	-0.00546	0.00650*
	(0.00433)	(0.00261)	(0.00269)	(0.00425)	(0.00606)	(0.00361)
Manuf. v.add.	-0.000860	-0.00295*	-0.00444**	0.00489	0.00390	-0.00442
	(0.00311)	(0.00156)	(0.00172)	(0.00465)	(0.00767)	(0.00307)
Constant	4.724***	3.589***	3.978***	-1.154***	3.853***	3.459***
	(0.289)	(0.0980)	(0.118)	(0.298)	(0.521)	(0.218)
WIID dummies	YES	-	-	-	-	-
Year FE	YES	YES	YES	YES	YES	YES
Observations	598	1,827	1,805	120	205	486
R-squared	0.790	0.067	0.108	0.748	0.186	0.106
# of countries	72	110	110	33	27	91

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. The dependent variable is the logged Gini coefficient from the data source indicated in the top row. Swiid_n and swiid_m refer to net and market inequality, respectively. Silc denotes the EU SILC data and wb the WDI Gini coefficients. The dummies for the underlying WIID categories are included in column 2 but are not shown to save space (available upon request).

Figure 3.A.1: Development of the Theil index in the countries of the MENA region

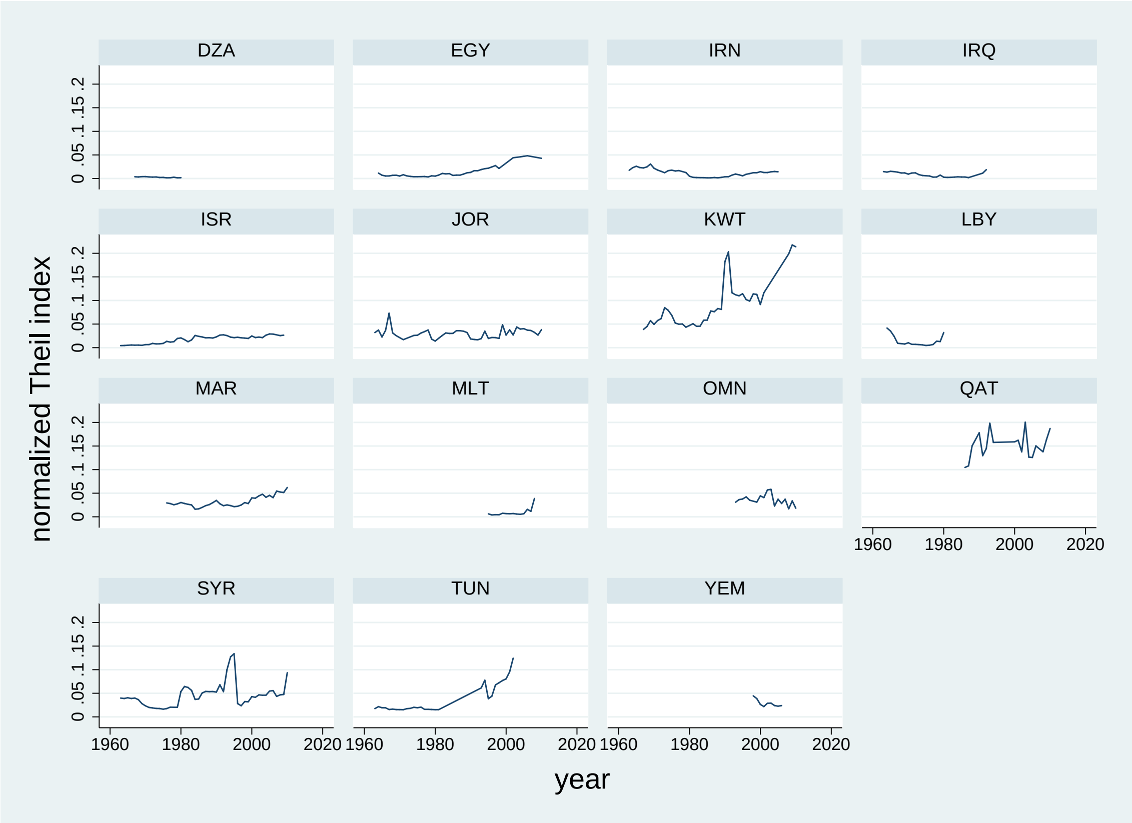

Notes. The years 1982-1988 in Tunisia rely on linearly imputed values. This means that the increase in inequality from the low level in the early years until 1981 to the peak in 1989 can, theoretically, occur less continuously - and not necessarily in a monotonous manner in any of the imputed years. Within Tunisia, the spike in 1989 is attributable to huge increases in the wage bills in several sectors, most notably, 15 and 18.

Figure 3.A.2: Size of the within-component and sectoral coverage by country

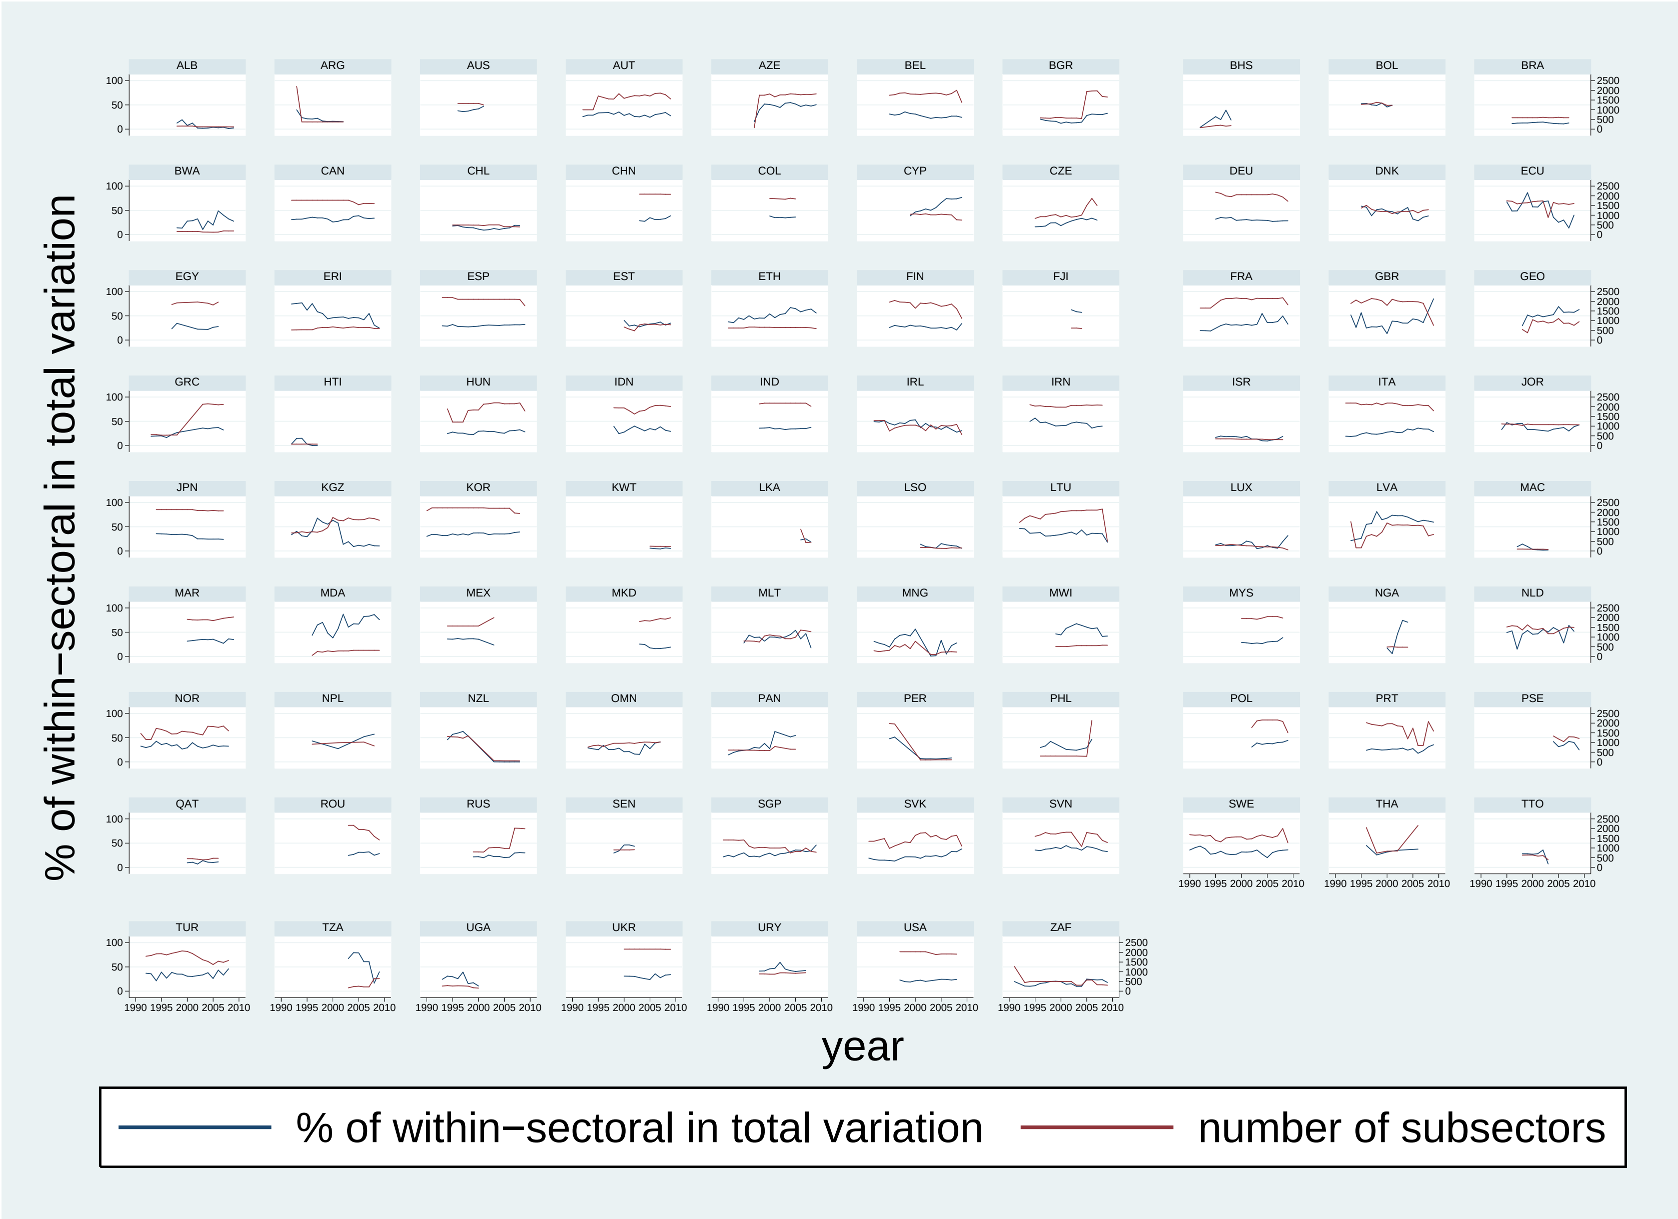

Figure 3.A.3: Log-normality of the (normalized) Theil index

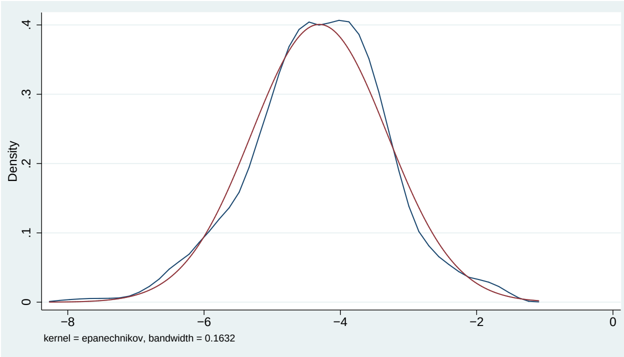

Figure 3.A.4: Kernel densities of the log differences between the Theil index and Gini coefficients of inequality

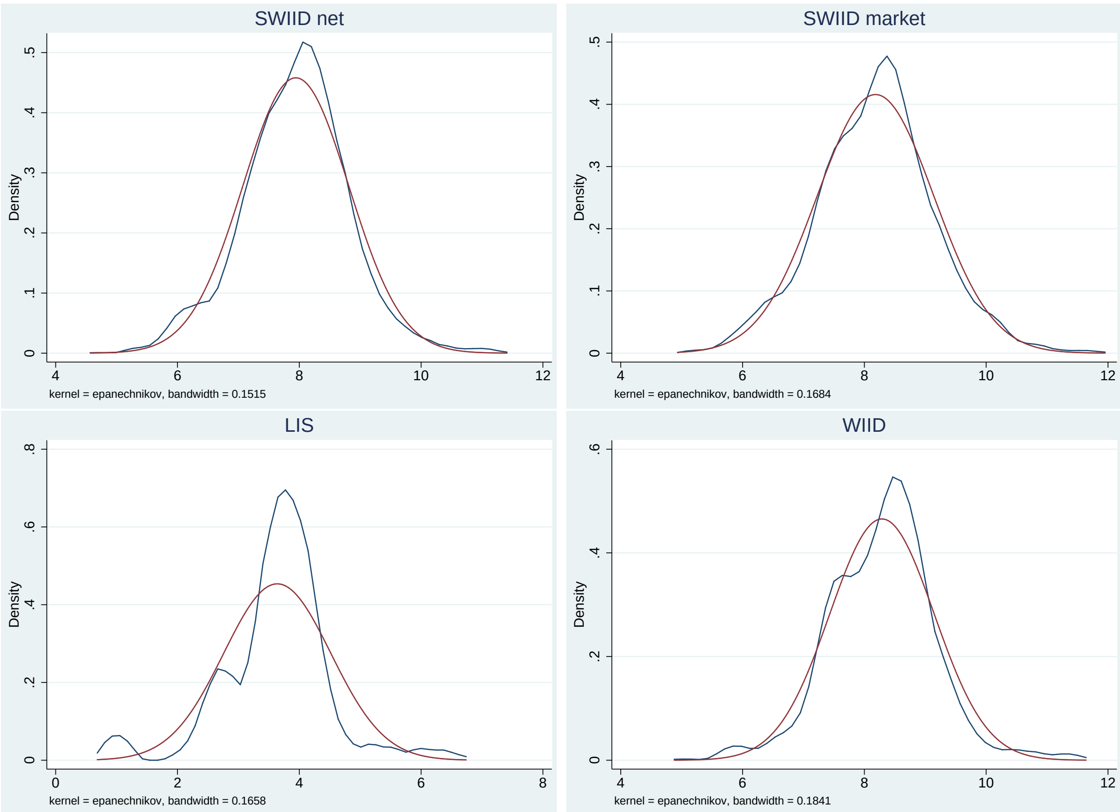

3.B Appendix

Imputation using fitted values from linear regressions

Filling missing values with predictions obtained from a regression permits the exploitation of further available information provided in the UNIDO industrial statistics. For the index of inter-industry wage inequality, only the data on wages and employment are needed. However, the UNIDO dataset also contains information on other sector-level characteristics. For the prediction of missing values, I use data on the number of establishments and on output as additional explanatory variables. If only one of the two variables needed for the computation of the index is missing, the other one is used in the regression as well. A time trend is also included in the set of potential regressors.

Once the first set of fitted values has been obtained from all available regressors, the next step is to assess the plausibility of the obtained prediction of the missing. Checking for plausibility is crucial for two reasons: First, the Theil index uses logarithmic transformations, which does not allow the inclusion of negative values. Second, because the index is based on the *ratio* of shares, a too-large or too-small number in one variable has the potential to affect the sector's contribution rather substantially and lead to disruptions in the series which may be unwarranted. In other words, the aim of the imputation is to arrive at plausible values for the missings, which at the same time should keep the series of inequality statistics smooth.

A good fit of the surrounding data points provides some indication of the appropriateness of the underlying regression model for a particular missing value. While the R^2 seems like an obvious candidate to judge the general goodness of fit, a high R^2 can sometimes be misleading, especially when the time series is long. In several cases, the fit is very good for part of the data, but captures relatively little of the variation in other parts. Whether the fitted values are useful for imputation then depends on where the missings are located. Each and every fitted value is therefore checked individually, and the regression is adjusted if necessary.

If an imputed value is deemed implausible, there are 3 principal ways to adjust the regression: (1) changing the regressors, (2) changing the time period, or (3) changing the imputation method. Only the first two options are discussed in the following, whereas the other imputation methods adopted are presented in the next subsections.

The first step is, of course, to identify "bad" fitted values, and predictions which are considered implausible for reasons other than a generally poor fit of the surrounding data points. Deviations of more than 30 percent of the fitted- from the actual values of the data points surrounding the missing are considered problematic and warrant changes in the regression model used. Then there are problems which arise occasionally despite a relatively good fit. The most obvious is a negative fitted value, which is conceptually impossible for wages or employee numbers. Along the same lines, even if the overall fit is good and the predicted value is positive, it can still be implausibly low or high. This basically happens when the values in the forcing variables suggest a value very different from the one obtained, and is mostly caused by large changes in one of the predictors to a level which does not occur elsewhere in the underlying data. Similarly, predicted values can be very different from their "surrounding" values and this is clearly not warranted by an extreme value in one of the forcing variables. These problems are of course related in many cases. In particular, negative values are just special case of an implausibly low fitted value. Similarly, their causes as well as the strategies for addressing them apply for several of the above cases.

Once a problematic fitted value has been identified, the next step is to check the coefficients of the individual variables to see whether a single regressor is driving the result.[51] Things that may indicate problems are negative coefficients (given that the initial reason for including the regressor was the assumption of a positive relationship) or a very large (or very small) size of an individual coefficient. Often, dropping the respective variable - which can also be the constant - solves the problem and yields a more realistic estimate of the missing value. However, it is not always possible to clearly identify an individual variable causing the problem. In many cases, all variables are useful in predicting a missing value, and it is not the set of variables but the time period which needs to be changed. This is especially true for long time series since the association between some of the variables is likely to not remain constant over a time span of 30 years or more, and sometimes changes visibly already in shorter time periods.

Some examples and illustrations with graphs and tables of the underlying data will help demonstrate the conjunctures encountered. Starting with the case of negative fitted values, suppressing the constant forces the regression line through the origin - an all but reasonable assumption, which helps to resolve the problem in many instances. An example is given below for missings in wages in Bolivia in sector 34 between 1971 and 1973, illustrated in table 3.B.1.[52]

As mentioned previously, implausible fitted values can be driven by outliers in one of the regressors. An example is the case of El Salvador shown in table 3.B.2 below, where the regression yields a very small number for the missing in wages in 1992 in sector 34. This is clearly due to the very low value of 22 of the explanatory variable "employees" in that year in comparison to the rest of the data for this sector, where employee numbers are always above 100. Obviously, the resulting wage number should also be substantially smaller than before, but is arguably not in the 3-digit range, as indicated by the value preceding the missing which is still around one third of the larger values in the later years. Here, suppressing the constant alone does not solve the problem. Only when the year 1998, containing substantially larger numbers for both wages and employees, is also omitted from the regression does it yield plausible numbers for the missing wages. In this case, plausibility is not only assessed through the deviations of the fitted values for the surrounding values, but also from observations with similar values for other regressors (in this case, establishments).

[51] While significance may seem like an obvious indicator of whether or not a regressor is useful, in many cases, the number of observations is too low to allow a judgement of which regressors to keep based on or the significance of the estimated coefficients - and whether or not they are robustly related to the regressand during the time period of interest.

[52] An alternative would be to allow for a different functional form, e.g., by including a cubic term. However, due to the often few degrees of freedom, this is not always feasible. Given the theoretically valid assumption of a constant of zero, this approach is hence preferred.

Table 3.B.1: Example of Bolivia: suppressing the constant

					without constant		with constant	
Year	Empl.	Wages	Estbl.	Output	Fv	%dev	Fv	%dev
1970	59			166667	3588		**-112731**	
1971	65			166667	37132		**-102628**	
1972	25							
1973	33			50000	13956		-166461	
1974	166	100000		550000	112416	12	100131	0
1975	150	150000		750000	137362	8	90243	40
1976	318	750000		3900000	617896	18	641739	14
1977	352	650000		4300000	681558	5	733099	13
1978	313	550000		4800000	744133	3	710063	29
1979	478	1400000		8750000	1337179	4	1324724	5
1980	513	1200000		10280000	1560859	30	1514124	26
1981	577	1680000	22	15960000	2377494	4	2106231	25
1982	363	562500	14	4203125	670153	19	743361	32

Notes. Fv is short for fitted values. %dev is the deviation of the fitted from the observed values in %.

Table 3.B.2: Example of El Salvador (sector 34): suppressing the constant

					without constant				with constant	
					without 1998		with 1998			
Year	Empl.	Wages	Estbl.	Output	Fv	%dev	Fv	%dev	Fv	%dev
1991		124740	1	1995841						
1992	22		1		96187		773		234452	
1993	135	306463	5	3196204	242767	21	624579	104	413323	35
1994	112	520464	2	3403122	212973	59	539949	4	391284	25
1995	152	309781	16	259864	264881	14	783491	153	462313	49
1996	273	731696	10	4286122	421835	42	1448971	98	653002	11
1997	223	571249	11	3397116	357026	38	1223695	114	591077	3
1998	968	7154426	8	20503255	1323221	82	5139651	28	1703586	76

Notes. Fv is short for fitted values. %dev is the deviation of the fitted from the observed values in %.

A generally bad fit is frequently caused by a particular data structure, wherein one can observe 2 different "regimes" in the development of wages and employees over time. Pooling these together in one regression yields a mediocre fit for both regimes. Excluding the years which display a different pattern from the one where the missings are located often solves the problem. Again using the example of El Salvador, in the first spell of data (pre-1985), wages are fairly stable, whereas in the second data spell (post-1990), they are much more dispersed and display high growth rates. Only post-1990 values are therefore used for the imputation of wages in 1992 in sectors 26, 27, 28, and 31. To give a better impression of this type of data structure, wages in these sectors are plotted in figure 3.B.1.

Figure 3.B.1: Log-normality of the (normalized) Theil index (wages in mn. USD)

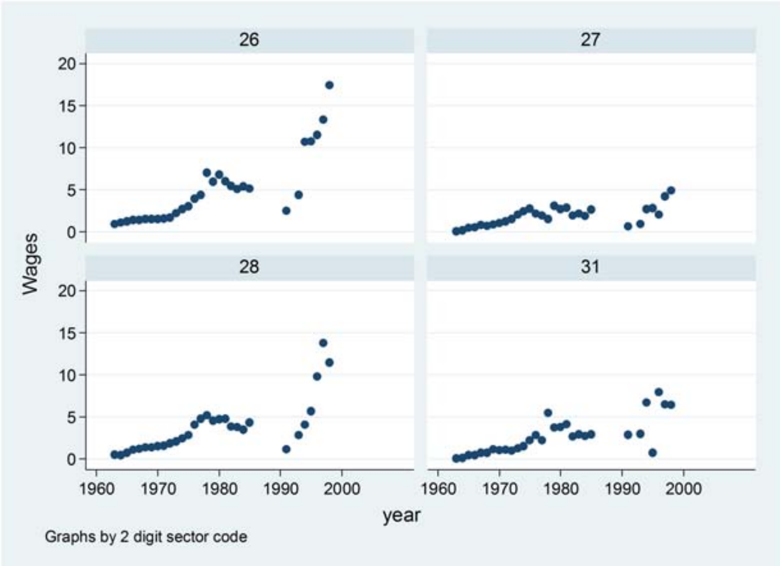

That the association with a forcing variable is deficient - as indicated most clearly by a large negative coefficient - occurs repeatedly in the estimations. It is easy to detect, and the straightforward thing to do is drop the respective variable. This improves the result in most instances. An example is Mozambique, where the "establishments" variable has a negative and relatively large coefficient in sector 34 for explaining employee numbers, and produces a correspondingly poor result with partly negative fitted values. The exclusion of the variable leads to a substantial improvement of the fit and yields positive values. Table 3.B.3 contains the raw data, and table 3.B.4 displays the results with and without the exclusion of establishments.[53]

[53] Note that although the "output" variable also has a negative coefficient, its effect is much smaller and its exclusion does not lead to a better fit, nor does it solve the problem of negative fitted values. Also note that the variable turns positive once establishments have been excluded from the regression equation.

Table 3.B.3: Example of Mozambique (sector 34, 1997/98 employees): dropping a variable

					without estbl.		with estbl.		without output	
Year	Empl.	Wages	Estbl.	Output	Fv	%dev	Fv	%dev	Fv	%dev
1986	1240		5	7892849	1433199		4335966		1179593	
1987	2710			7160227	1317665					
1988	2920		10						980049	
1989	2760		9						916807	
1990	2487	1767323	9						843129	52
1991	2276	943206	9	6923114	909802	4	1220781	29	76946	18
1992	1102	590888	11	3145974	732922	24	600248	2	674896	14
1993	1138	640900	11	2064663	610363	5	613682	4	601218	6
1994	1049	875039	10	2618161	520744	40	888217	2	537976	39
1995	446	179071	11	2179552	411135	130	196929	10	45386	153
1996	911	327082	10	3541943	337816	3	313905	4	390618	19
1997			**24**	**4005683**	**246389**		**-8087087**		**170818**	
1998			**35**	**17402439**	**415559**		**-1.7E+07**		**-17670**	
1999			19						75647	
2000	475	68430	17						22842	67

Notes. Fv is short for fitted values. %dev is the deviation of the fitted from the observed values in %.

Table 3.B.4: Example of Mozambique (sector 34, 1997/98 employees): regression output

Dep. var.: Wages	(1)	(2)	(3)
Output	0.0201	-0.195	
	(0.258)	(0.0366)	
Establishments		-579,534*	-10,437
		(50,212)	(72,769)
Year	-100,771	-197,187**	-73,679
	(102,847)	(15,080)	(68,392)
Constant	2.014e+08	4.004e+08**	1.476e+08
	(2.049e+08)	(3.038e+07)	(1.358e+08)
Observations	5	5	6
R^2	0.329	0.995	0.567

Notes. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. The three columns contain the regression output for the three sets of fitted values/ %-deviations of table 3.B.3.

Individual variables can also cause a generally bad fit and dropping the variable often helps resolve the problem. An example is sector 36 in Bulgaria, where the inclusion of the variable "establishments" leads to large deviations and negative values in some years (see table 3.B.5). Here, the large coefficient on the variable (shown in table 3.B.6) is indicative of the problem. This example also demonstrates how a short time period restricts the options for achieving a better fit: excluding the later years containing substantially higher numbers for establishments would theoretically also be possible, but would leave even fewer years for estimation. Excluding the "establishments" variable is therefore preferable in this context.

Table 3.B.5: Example of Bulgaria (sector 37, 1996 wages): dropping a variable

					without estbl.		with estbl.	
Year	Empl.	Wages	Estbl.	Output	Fv	%dev	Fv	%dev
1996			26	736414	95838.74		771319.9	
1997	100	117882	26	1413299	103559.1	**12**	509868.9	**333**
1998	100	120108	25	1411643	102262.8	**15**	187921.1	**56**
1999	123	131675	18	2722748	118410.9	10	-433161.2	429
2000	303	370117	19					
2001	290	219709	24					
2002	426	385656	40					
2003	550	772205	46	47902118	713671.8	8	493296	36
2004	158	206335	57	5079010	143350.2	31	292953.3	42
2005	282	387516	64	7623259	175884.7	55	420798.6	9
2006	1211	2174736	62	2.53E+08	3431034	58	2952456	36
2007	1540	5682796	86	3.56E+08	4805079	15	5169121	9

Notes. Fv is short for fitted values. %dev is the deviation of the fitted from the observed values in %.

Table 3.B.6: Example of Bulgaria (sector 37, 1996 wages): regression output

Dep. var.: Wages	(1)	(2)
Output	0.0133***	0.0119***
	(0.00260)	(0.00241)
Establishments		52450
		(31,883)
Year	-1274	-269478
	(94,069)	(182,145)
Constant	2.630e+06	5.373e+08
	(1.882e+08)	(3.634e+08)
Observations	8	8
R^2	0.909	0.946

Notes. Standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. The two columns contain the regression output for the two sets of fitted values %-deviations of table 3.B.5.

There are also reasons to ex ante exclude available regressors from the estimation. This often concerns the "establishments" variable, where the absolute numbers are sometimes very low (single-digits) and/or where there is little variation (e.g., in Lesotho, the number of establishments is constant and below 10 in several sectors for the 2001 to 2008 period and varies only by 1 in a few others). Another limitation to the inclusion of regressors is the number of observations. Due to the unbalancedness of the data, including an additional variable often reduces the years available for estimation. Again using the example of Lesotho, for predicting missing employee numbers in sectors 17 and 24 in 1980/81, the "output" variable is available only for years post-1981, but not for the earlier ones. As its inclusion would reduce the number of observations to a level where no degrees of freedom are left for estimation, it is dropped for the estimation. Omitting the constant also increases the degrees of freedom and performs better in other sectors where the output variable is more important in predicting the missing and is therefore retained.

Implausible values of any kind can of course also arise from a simple outlier in one of the forcing variables, in which case it suffices to exclude the respective year from the regression. An example from Puerto Rico is shown in table 3.B.7, wherein employee numbers drop to 650 in one year from a level of around 4000 in all other years, worsening the fit substantially.

Table 3.B.7: Example of Puerto Rico (sector 26, 1987/88 wages): outlier years

					with 1998		without 1998	
Year	Empl.	Wages	Estbl.	Output	Fv wages	%dev	Fv wages	%dev
1987	4520		162		2.33E+07		4.09E+07	
1988	4630		164		2.91E+07		4.87E+07	
1989	4950	34800000	167	5.00E+08	3.11E+07	11	5.75E+07	65
1990	3920	36900000	159	5.01E+08	5.81E+07	57	6.05E+07	64
1991	3620	92500000	156	5.02E+08	7.17E+07	23	6.66E+07	28
1992	3460	92100000	41	5.06E+08	9.42E+07	2	8.10E+07	12
1993	3370	93800000	151	5.25E+08	9.25E+07	1	8.06E+07	14
1994	3600	98300000	157	5.67E+08	9.57E+07	3	8.87E+07	10
1995	3620	1.03E+08	162	5.42E+08	1.03E+08	0	9.59E+07	7
1996	3610	1.06E+08	196	6.32E+08	1.07E+08	2	1.01E+08	4
1997	3760	1.09E+08	184	7.00E+08	1.14E+08	5	1.10E+08	1
1998	**650**	96400000	**37**	7.03E+08	1.92E+08	99	1.12E+08	16
1999	4010	1.19E+08	214	7.69E+08	1.22E+08	3	1.24E+08	4
2000	4340	1.31E+08	197	7.38E+08	1.26E+08	4	1.34E+08	2

Notes. Fv is short for fitted values. %dev is the deviation of the fitted from the observed values in %.

If none of the described regression-based solutions yield any useful results, another type of imputation is applied. The idea is similar to a simple linear interpolation but still exploits some of the information contained in other variables. The method assumes a constant co-movement of the missing with another variable (that is, not using the "year" variable, which would be the case in a "normal" linear interpolation) and traces its development in the missing years. Missing employees in 1995 in sector 35 in Slovenia provide an example, shown in table 3.B.6. Here, the numbers rise so drastically for other variables in the years following the missing that no regression model can be found which results in acceptably low, but still positive, numbers. In this case, the problem stems from a lack of support for numbers of this magnitude in the data. Employee numbers have to be imputed differently and are assumed to move in accordance with wages[54] in the concerned years.[55]

[54]The imputed values for wages are used, which have a fairly good fit in the early years (shown in the last two columns table 3.B.6).

[55]The relatively high correlation of 0.78 between the two variables supports this assumption.

Table 3.B.8: Example of Slovenia (sector 35, 1995 employees): linear interpolation

Year	Empl.	Wages	Estbl.	Output	Fv empl.	Fv wages	%dev wages 1997	%dev wages 1996
1995			33	84922771	**1373**	17911636		
1996	144		33	53110129	144	7478213		6
1997	143		43	39017597	143	2982585	4	25
1998	835	10097149	51	57697994	835	10097149	4	15
1999	2808	39551278	60	1.27E+08	2808	39551278	5	19
2000	2818	34440932	64	1.25E+08	2818	34440932	6	8
2001	1362	12833498	69	69103452	1362	12833498	5	3
2002	1373	17954450	87	1.02E+08	1373	17954450	3	32
2003	1317	20826801	94	1.16E+08	1317	20826801	11	36
2004	1332	9965221	119	67265244	1332	9965221	25	19
2005	2027	39793868	128	1.44E+08	2027	39793868	8	4
2006	2568	51433489	157	2.06E+08	2568	51433489	4	15
2007	2587	60224473	170	2.55E+08	2587	60224473	31	26
2008	1496	42479950	86	1.93E+08	1496	42479950	4	29
2009			84			17911636	4	
2010			81			7478213	5	

Notes. Fv is short for fitted values. %dev is the deviation of the fitted from the observed values in %. The estimation for wages in 1997 is based on a regression of wages on employees, output, establishments, and a time trend, whereas the 1996 value is based on output only.

Other imputation approaches

If no information is provided for other variables which would allow a regression-based imputation, a simple linear interpolation between the surrounding values is performed. This is equivalent to the above approach of imputation alongside another variable, but always using the "year" variable. For example, in the case of Bangladesh, there is no data in the years 1993 and 1994 but values for both wages and employees are available in 1992 and 1995. In sector 15, the number of employees is 107882 in 1992 and 126220 in 1995. The resulting values for 1993 and 1994 are calculated as (126220-107882)/3+107882 $\approx$ 113995, and (126220-107882)/3*2+107882 $\approx$ 120107. Imputation based on linear interpolation hence always implicitly assumes a linear development over time of the target variable for the missing years between the two surrounding data points.[56]

A disadvantage of the linear interpolation approach is that is requires both a start- and an end-observation for the missing time period. If a missing is located in the first or last year of the available data, the method cannot be applied. Instead, for missings located at the beginning or the end of a data spell, a time trend is used to extrapolate values when no other information is provided by the dataset - again exploiting the year variable.[57] The

[56]In a few cases, means imputation is based on starting or ending values which are the result of a regression-based fitted value imputation (e.g., sector 27 in Fiji for the missings in wages in 1994-95, which use as the "starting" observation the fitted value of 1993). This is done because the alternative would be to start linear interpolation at the closest data points provided by the raw data, which means that the fitted value would be overruled by the linearly interpolated one. This runs contrary to the initial idea that regression-based fitted values are always favoured over linear interpolation, as they incorporate all of the information available from the data.

[57]One could also use a time trend to impute values missing "in the middle" as an alternative to the simple linear interpolation described above. Using a time trend is advantageous when there are outliers at the beginning or end of a linear imputation, as well as a discernible time trend in the data. Otherwise, it is less

same procedure as with the regression-based approach is applied, including the option to drop the constant or change the time period when the fit is bad.

Again, there are several cases where it is not possible to find a good fit which would support extrapolation based on a time trend. In those cases, the first (or last) available value - which is in some cases an imputed one[58] - is then repeated in the missing years. This has, e.g., been done in Fiji in 1995, where the 1996 value is used to fill the missing in sector 20. Whether such a procedure is reasonable also depends on the development of the data in the preceding or following years: if they have been relatively stable, using the same values seems valid.[59]

Of course, the discrete procedure of imputing data on a case-by-case basis is inherently arbitrary. This applies not only to the imputation procedure, but also to the preceding decision of whether or not to impute in the first place. As a rule of thumb, no sector is used in the final index in which more than 50% of the data need to be imputed.

suppositional to assume that values do not range outside the value of the start and the end year of the gap. Given that the goal is to tamper with the data as little as possible, if no further information is provided in the data which would point towards the missing going into a particular direction, it is desirable to merely preserve the ratio of wage and employee numbers in order to maintain time coverage, but influence the inequality index as little as possible. Linear interpolation effectively means that the resulting contribution of the imputed missing will lie between that of the start and that of the end year, and is therefore the least intrusive option and preferred over the use of a time trend.

[58] An example for when an imputed value was used to fill missings in the first few missing years (1970-1973) is sector 27 in Indonesia.

[59] There are only two cases with imputation approaches different from the ones described, but based on the same techniques. The first one is Bulgaria, which is the only case where a squared term has been employed to impute employee numbers post-2003 due to the clearly discernible inverse U-shaped development of employee numbers in sector 23. The second case is Tunisia, where the fitted values for sectors 22, 28, and 36 are located in a time period relatively far from another data spell containing support for both variables involved in the imputation. The fitted values are the result of an average of two imputation approaches with very different results, but an equally good fit in their respective (non-overlapping) parts of the data. The resulting fitted values line up nicely with the values of the time series following the missing.

Chapter 4

The Impact of Trade on Wage Inequality in Developing Countries: Technology vs. Comparative Advantage

4.1 Introduction

In the 1980s, developing countries considerably lowered barriers to international trade, thereby substantially boosting trade flows. This comprehensive economic change has not been without distributional consequences. The Heckscher-Ohlin (H-O) theory (Heckscher 1991) yields clear predictions of the effects of trade on the distribution of income among production factors. Their relative abundance is also the source of comparative advantage in international trade and countries abundant in one production factor will specialize in the production of goods relatively intensive in that factor. The relatively abundant factor will gain, while the scarce factor, experiencing the opposite effects, will lose from trade (Stolper and Samuelson 1941).

Developing countries, relatively abundant in low-skilled labor, would hence specialize in low-skilled labor-intensive production. Because low-skilled labor is generally located at the lower end of the wage distribution while high-skilled labor forms the upper end, wage inequality should decrease in developing countries as a result of increased exposure to international trade. Furthermore, because capital is complementary to high-skilled labor in many cases and relatively scarce in developing countries, the same should be true for income inequality (Krusell et al. 2000, Goldin and Katz 1998).

Available data on both wage and income inequality describe a reality very different from what one would expect based on traditional trade theory after the large increases in world trade volumes. Inequality has been rising not only in the industrialized countries

[*] **Acknowledgements:** I would like to thank Stephan Klasen, Axel Dreher, Sunil Kanwar, Inmaculada Martínez, Bernhard Brümmer, Atika Pasha, and the participants of seminars and conferences at the ISI Delhi, the Delhi School of Economics, the Courant Centre Göttingen, the AEL workshop, and the RTG Globalization and Development workshops for helpful comments and suggestions, and Mathis Richtmann for valuable research assistance.

but also across the developing world. The correlation between the expansion of world trade and rising inequality does, of course, not imply causality. There are many factors related to both globalization and trade which may conflate or counteract any equalizing effects of trade on the income distribution.

Several papers have shown that trade has a differential impact on inequality in high- and low-income developing countries and that this effect differs depending on the trading partner as well (e.g., Gourdon 2011, Meschi and Vivarelli 2009). The differential impact has been attributed to technology transfers from rich to poor countries, although this transmission channel is rarely tested directly (one notable exception being Conte and Vivarelli (2011), who find evidence of such transfers). Rising skill premia have indeed been shown to increase wage inequality not only in developed countries, but in developing countries as well (for an overview of the literature, see Vivarelli 2014). Failing to account for the source of this development leaves open the question of whether technological change does in fact arise through trade, or whether it could be domestic technological change stemming from technological innovation within the respective country itself which raises skilled wages. Taking technological change into account is important because it is potentially driving both exports and wages in certain sectors and may thereby introduce a spurious correlation between trade and wage inequality. Most studies "assume away" domestically induced technological change in developing countries and argue that all technological advancements stem from external sources. To support their claim, they refer to the low level of research and development activities, as first stated by Coe et al. (1997). While it may be true that there is very little domestically induced technological advance in earlier time periods (before the 1990s) for certain countries, it does not seem plausible for upper-middle income countries such as South Korea, Spain, or Slovenia even in earlier years of the sample periods used, or for countries like India in the early 2000s.

Another shortcoming of many empirical papers using the H-O model to test the effects of trade on the income distribution is the timing of the trade effect. Certain factors of production cannot be assumed to be mobile between sectors in the short run, and hence the predicted effects may not be visible in a contemporary or one-year lagged specification. The Ricardo-Viner model (Viner 1932, Jones 1971, Mussa 1974), introducing specific factors into the theoretical framework, is often interpreted as a short-run version of the H-O model and is thus likely to better capture the effects within the time horizon of a few years which can be feasibly estimated with the available data. According to the model, immobile factors of production may lose from trade if employed in import-competing industries, even if they are overall relatively abundant. The opposite holds true for relatively scarce factors which are employed in exporting industries and which may gain from trade in the short run. With skilled labor being one of the most frequently cited examples of a specific factor, the model is highly relevant in the context of this paper in which skill premia are a key mechanism in driving up wage inequality.

This paper addresses the identified shortcomings of previous studies in several ways. First, it directly measures the technology content embedded in trade by categorizing trade flows into different technology levels. Second, it includes a new measure of technological change to address potential omitted variable bias. The measure captures movements in the technological frontier, which is estimated using data envelopment analysis (DEA) and based on the same raw data used in the inequality index. It is hence able to control for advancements in technology in exactly those sectors included in the inequality measure as well. Differentiating between imports and exports helps to disentangle the two technology transmission channels. Furthermore, different types of hypotheses can be tested on the two variables. In particular, the H-O model does not provide any insights into the effect of imports, whereas the specific factors model provides clear distributional implications with respect to import-competing industries.

In order to maximize the time coverage, a Theil index of between-sectoral wage inequality covering the years 1970-2010 has been constructed (see chapter 3). It is based on the UNIDO industrial statistics, covering manufacturing industries in a large number of developing countries. A major advantage of the lengthy time coverage with a maximum of 41 years is that fixed effects estimation delivers reliable estimates in terms of Nickell bias despite the dynamic specification of the econometric panel data model. The sample for the preferred specification contains 58 developing countries over an average time span of 16 years, and results from a GMM estimation confirm that Nickell bias does not affect the estimates.

Results suggest that while technology transfer through trade does play a role in driving up wage inequality in developing countries, it is important to control for domestic technological change as some of the effects otherwise attributed to trade disappear once the measure is included. Trade does not seem to per se drive up wage inequality, and the traded technology appears to play an important role. The disequalizing effects stem primarily from trade in medium-technology intensive goods and occur mainly in countries in the second education quartile, that is, with medium skill endowments. Few results are found for trade with high-technology goods, which casts doubt on the hypothesis that it is technology transfer in these goods causing the disequalizing impact of trade with developed countries in developing countries found in previous studies.

Although there is a large recent literature emphasizing the impact of trade on wage inequality within industries and occupations (broadly categorized into effects due to heterogeneous firms, labor market frictions, and incomplete contracts), the present study focuses exclusively on inequality between sectors and refers to sector-based classical trade theory. Although this choice is partly dictated by the nature of both the wage inequality data as well as the sector-based classification of the trade data, both of which are only available at a sectoral level, this paper addresses several of the previously identified shortcomings of existing sector-based studies. Results suggest that the lack of strong results pertaining to the effects of the neoclassical, sector-based mechanisms is at least partly due to flaws in the empirical approach of estimating the relationship between trade and inequality. Moreover, some of the explanations subsumed as challenges to the neoclassical trade theory, such as Feenstra and Hanson's (1996) offshoring argument (as, e.g., in Harrison et al. 2011) can be easily incorporated into the sector-based model.

Taking a detailed look at the available inequality data, several studies have identified changes in the upper quintile of the income distribution to be the main driver of inequality. The income share of the upper quintile increased at the expense of the middle part of the distribution while there has been little change at the bottom (e.g., Jaumotte et al. 2013). Goldberg and Pavcnik (2007) find a pervasive increase in skill premia across developing countries during the 1980s and 1990s, which translates in most cases into an increase in wage inequality. The two decades are particularly interesting and are the focus of most empirical studies because not only has inequality gone up, but many countries have opened their economies to trade at the same time.

The determinants of the increase in income and wage inequality in advanced economies are relatively well explored. Even though the co-movement of trade and inequality is in line with the H-O/S-S predictions, trade has been found to be only of minor importance for the large increases in inequality in the 1980s and 1990s. Rather, skill-biased technological change (SBTC) has been identified as the main cause for the changes in the distribution of wages and incomes (e.g., Berman et al. 1998; see Card and DiNardo 2002 for a more critical review of the SBTC hypothesis and Kurokawa 2014 for a survey of trade-based versus other explanations). The basic reasoning behind this is that technological progress is complementary to high-skilled labor and consequently raises demand for the highly skilled (Acemoglu 2003). There is evidence that SBTC is present in developing countries as well, and that trade introduces or reinforces SBTC in those countries (Berman and Machin 2000, Conte and Vivarelli 2007). More recent studies focusing on European countries ascribe a larger role to trade in increasing inequality through exporter wage premia (Klein et al. 2013; Egger et al. 2013; Baumgarten 2013), which is even more pronounced if import penetration is also accounted for (Du Caju et al. 2012). The latter study even finds that the negative impact of imports on wage levels is larger for trade with developing countries. However, it remains unclear how large the contribution of trade is to overall wage inequality, and whether corresponding effects are present in developing countries.

The geographical distribution of trends in income inequality points toward another explanation, which is complementary to the SBTC hypothesis. While the advanced and newly industrializing countries in Asia, Latin America, and Europe have experienced increasing income inequality over the 1980s and 1990s, this is not generally true for low-income countries, particularly in Sub-Saharan Africa (Jaumotte et al. 2013). Several countries, in particular in Latin America, have also experienced marked decreases in income inequality since the mid-1990s (Cornia 2014). This differentiated pattern of development of income inequality across countries lends support to an argument first introduced by Wood (1997), which explains the apparent lack of an equalizing effect of trade by making a more detailed distinction between country groups. Trade between developing countries, often labeled "South-South trade," obviously does not fit in with the dichotomy of "North-South" trading partners and their relative endowments assumed in most H-O-based models. What constitutes a comparative advantage in trade between "Southern" countries must be established before any predictions about the effect of trade on inequality between developing countries can be derived.

In the following, the theory behind the technology and the South-South trade hypotheses will be explained in more detail. Empirical evidence on the roles of trade, technology and South-South trade as well as the effects of their interrelations on income inequality will be reviewed thereafter. The empirical analysis is covered in section 4.3, which introduces the data and motivates the empirical specification. Estimation results are discussed in section 4.4. Robustness checks are presented in section 4.5, and section 4.6 concludes.

4.2 Literature review

4.2.1 (Skill-biased) technological change

Katz and Autor (1998) and Conte and Vivarelli (2011) summarize the various patterns on the production side of the economy indicating the occurrence of SBTC. Among them is the constant or increasing ratio of high-skilled to low-skilled workers despite rising skill premia, and thus relative wages, for the highly skilled. This phenomenon has recently also been observed in several developing countries (e.g., Berman et al. 1998, particularly in emerging economies such as India, Hong Kong, and several Latin American countries (for a review see Goldberg and Pavcnik 2007). Berman and Machin (2000) find evidence of SBTC, measured by the share of non-production relative to production workers, in middle-income, but not in low-income developing countries. They also notice that the same industries are affected by SBTC in OECD and in developing countries and infer that SBTC in developing countries is driven by a transfer of technology from industrialized countries. Trade is an obvious candidate as one of the vehicles of technology transfer. It can act as a catalyst of (skill-biased) technological change[1] in developing countries, thereby reinforcing the disequalizing effect of rising skill premia. Imports may provide formerly unavailable goods that embody new technology complementary to skilled labor. They can also be investment goods that enable the introduction or modernization of production processes (Pissarides 1997), or final goods that allow for reverse engineering (Meschi and Vivarelli 2009). Imported capital goods can also be substitutes for low-skilled labor and introduce labor-saving technology, which leads to a widening wage gap through the depression of low-skill wages (Behrman et al. 2007). Summarizing the above arguments as the ”import channel,” Meschi and Vivarelli (2009) also identify an ”export channel” through which SBTC is introduced in developing countries. Export partners in developed countries have certain demands on the quality and up-to-dateness of the products they import. They might therefore either directly assist their developing country partners in upgrading their technology and the skills of their workforce, or make an investment in such upgrading profitable. Intermediate goods can have effects through both the import- and the export channel. Feenstra and Hanson (1997, 2001) argue that their impact on wage inequality is particularly strong because demand for skilled labor does not only affect the exporting or export-competing industry, but also all the industries that use the intermediate goods as inputs, regardless of whether they trade the final product or not. They also point out that some industries are more suitable for outsourcing than others. Outsourcing is more present in industries in which the production process can be separated into more or less independent stages and in which the different steps of production entail

[1] The term ”skill-biased technological change” is in the original sense different from mere technological upgrading in developing countries, which is not necessarily skill-biased from a developed country point of view. However, since such upgrading frequently is skill-biased from the developing country’s perspective, the term will be used here to include both meanings.

large differences in the skill composition. Feenstra and Hanson (1997) find that these are mainly industries producing semi-durable consumer goods. The manufacturing sector therefore seems particularly prone to such effects.

Given the potential for technological catch-up, the effect of trade on technological upgrading may be particularly strong in developing countries, especially in emerging economies. Schiff and Wang (2004) show that developing countries benefit more from increased import volumes than developed countries in terms of productivity improvements. The adoption of new or upgraded technologies not only depends on their availability, but also on a country's capability to employ it and take advantage of it. If there is an insufficient supply of knowledge and qualified labor, or low domestic demand, new technologies will not be established. Acemoglu (2003) makes this point in his model of endogenous technological change: Technology used in developing countries prior to trade liberalization is adapted to local circumstances, thus complementing low-skilled labor. New technologies introduced via imports on the other hand are designed to match the mix of production factors in developed countries and are therefore skill-intensive from a developing country's point of view. The decision as well as the possibility to adopt skill-intensive technology depends on the ability of a country to use it and to benefit from it, which in turn depends on the composition of its labor force and the supply of skilled labor. Zhu's (2004) model relies on a similar assumption and introduces a link to the product cycle, wherein new, more skill-intensive goods developed in industrialized countries replace older ones. The production of the older goods is then transferred to developing countries and constitutes a new, relatively skill-intensive production technology there. As a consequence, skill premia rise in both country groups. Pissarides (1997) argues that even if a new technology is not skill-biased, its mere introduction requires skilled labor because new technologies have to be learned about and put into use. The effect on the demand for skilled labor is then transitory. This is also true if one considers that skill-biased technologies can be modified in a way such that they complement unskilled labor. This modification also requires a certain amount of knowledge and skilled labor. A similar point is made by Bernard and Jensen (1997) and Matsuyama (2007), who argue that the activity of exporting is skill-intensive in itself.

Given the above considerations, it stands to reason that an educational expansion fostering an increase in the supply of high-skilled workers is a prerequisite as well as an accelerator of SBTC in developing countries. At the same time, it depresses skill premia in the short run because of the time lag of new investments in more skill-intensive technology reacting to the increased abundance of skilled labor. Acemoglu (1998) finds evidence in the United States for both the short-run, equalizing effect of education on skill premia and the long-run effect, fostering skill-biased technological change and raising skill premia. In this paper, the short-run (supply) effect will be tested directly, whereas the long-run effect is implicitly incorporated into the classification of countries according to their relative skill levels.

4.2.2 South-South trade

The basic reasoning behind the South-South trade argument is that countries that are pooled in a rather undifferentiated manner under the label of "developing countries" are in fact so heterogeneous in terms of economic and human development that the relative abundance of production factors, and hence the impact of trade, differs vastly between them. While the unskilled workforce in the least developed countries generally benefits from trade because it can exploit its comparative advantage in low-skill production sectors, the case is different for middle-income countries, comprising also the newly industrializing countries. These countries have evolved to a stage where they no longer have a comparative advantage in unskilled labor. One can therefore not per se assume that trade with either developed or developing countries leads to a decrease in wage inequality in these countries. The fact that many developing countries felt the need to protect low-skill sectors through tariffs and other trade barriers prior to trade liberalization underpins the hypothesis that this is not where they had their comparative advantage. It rather shifted to medium-skill intense production, in particular when many developing countries with a large unskilled labor force - the most prominent example being China - entered the world market during the period of liberalization in the 1980s (Wood 1997).[2] The impact of trade with low-income countries in the low-skill, labor intensive sectors of middle income developing countries would then again be in line with the predictions of H-O/S-S: product prices fall and factor rewards are reduced - implying a larger wage gap. Davis (1996) has formalized this point in a theoretical model on the effects of trade liberalization on factor rewards within different groups of countries with similar endowments. It is hence crucial to differentiate between different kinds of developing countries in order to get clear results on the effects of trade on wages.

4.2.3 Empirical evidence

As previously mentioned, the results of "early" studies (meaning that neither technology nor trade between developing countries is taken into account) on the impacts of trade liberalization on the income distribution in developing countries are mixed. Most of them use the Gini coefficients from Deininger and Squire (1996) as their dependent variable, a few use quintile shares, and only one study analyses wage inequality. An unambiguously negative impact of trade on inequality is found by only few studies (examples include Bourguignon and Morrisson 1990, and Calderón and Chong 2001). Positive effects are identified by Lundberg and Squire (2003), Cornia and Kiiski (2001), and Spilimbergo et al. (1999). Barro (2000), Savvides (1998), and Milanovic and Squire (2007) all conclude that the disequalizing effects are stronger or only present in developing countries. Studies which find no effect at all include Edwards (1997), and Dollar and Kraay (2002, 2004) who find that average incomes and incomes of the poor are affected equally by trade.

Several authors have acknowledged the difficulty of drawing conclusions about the relationship between trade and income inequality from these studies because comparability

[2] Dollar and Kraay (2004) provide a list of developing countries they identify as "post-1980 globalizers" based on the increase in trade over GDP between 1980 and 2000 and backed by changes in tariff policies.

is limited due to differences in the countries and time periods covered, the choice of the inequality- and the openness variables, and the econometric specification and methodology used (Milanovic and Squire 2007, Lundberg and Squire 2003). Consequently, other approaches have been developed to explain the apparent lack of a clear-cut relationship between trade and inequality in developing countries, of which the SBTC and technology transfer arguments have received most attention. As for the South-South trade hypothesis, only two studies, by Gourdon (2011) and Meschi and Vivarelli (2009), explicitly incorporate trade between different groups of developing countries into their empirical analyses.

4.2.3.1 The role of technology: SBTC and technology transfer

A large number of country case studies investigate the interrelationships between technology, trade, and inequality in developing countries. Most of them find evidence for trade-induced technological change driving up skill premia and inequality - an exception being Ferreira et al. (2007), who conclude that trade has led to a decrease in inequality through sector reallocation effects of employment, as suggested by H-O theory. For a review, see Robbins (1996) on early evidence and Gourdon (2011) for more recent studies. The number of cross-country studies on the other hand is considerably lower. Zhu and Trefler (2005) find that wage inequality in developing countries in terms of relative wages of skilled to unskilled workers has increased due to trade-induced technological catch-up, measured by labor productivity. Zhu (2005) puts her theoretical model of technology transfer through product cycles to an empirical test in a panel of 28 US trading partners. Results indicate that product cycle trade leads to skill upgrading in countries which have a GDP per capita of at least 20 percent of the US GDP per capita, while no effect is found in the lower income countries. Conte and Vivarelli (2007) estimate the impact of "skill-enhancing technology import" from high income countries on the employment of the skilled and unskilled in low and middle income countries. According to their results, trade-induced technological upgrading entails not only a relative, but an absolute skill bias also since it not only increases the absolute employment of skilled workers, but decreases the number of unskilled workers as well. However, the analysis does not control for the supply of skilled and unskilled labor. Robbins (1996), including various direct measures of labor supply, finds that shifts in labor supply have large effects on relative wages, and concludes that labor markets are to some degree insulated from factor price equalization. López-Calva and Lustig (2010) argue that an educational expansion, lowering the gap between skilled and unskilled labor, is one of the main factors responsible for the decrease in labor income inequality observed in Latin America over the 2000s. This means that Conte and Vivarelli's (2007) results could suffer from omitted variable bias because the supply of skilled labor is not controlled for. In addition, not only imports but also exports can be a source of technology transfer. Finally, Jaumotte et al. (2013) measure technological change by the share of information and communications technology capital (ICT) in the total capital stock in their analysis of both advanced and developing countries and conclude that the main driver of inequality is technological change, above and beyond its effect through trade. Trade is found to reduce inequality, though mainly through exports of agricultural products, with no separate effect of manufactured goods.

4.2.3.2 Incorporating South-South trade

One of the two studies explicitly testing the South-South trade hypothesis while also taking SBTC into account is Gourdon (2011). To estimate trade-induced technological change, relative total factor productivity between skill-intensive and non-skill intensive sectors is regressed on North-South trade (between high-income and developing countries) and South-South trade (between middle-income and low-income developing countries) in a sample of 68 developing countries over 1976-2000. Inter-industry wage inequality is then regressed on North-South and South-South trade as well as the respective previously identified effects of technology transfer. This procedure allows for separately identifying the direct effect of North- and South-South trade on inequality and their respective indirect effect via technological change. Once technology transfer is controlled for, North-South trade has an equalizing effect on wage inequality while South-South trade increases inequality in both middle-income and low-income developing countries. While the effect in middle-income countries is direct, it operates through technology transfer from middle- to low-income developing countries in the latter. The analysis makes an interesting point in that trade-induced technological change in developing countries can originate not only from developed, but also from other developing countries.

Meschi and Vivarelli's (2009) analysis combines both the technology transfer and the South-South trade hypotheses in a sample of 65 developing countries from 1980 to 1999. The analysis relies on the UTIP-EHII measure of income inequality, which combines the Deininger and Squire (1996) dataset with the UTIP-UNIDO wage inequality data. Trade flows are decomposed by their origin and destination countries and it is found that trade from and to developed countries worsen the income distribution, while trade with other developing countries has an equalizing effect. Further results confirm the technology transfer hypothesis: trade with developed countries has a negative impact only in middle-income developing countries, while the effect in low-income countries is insignificant. Trade between low- and middle-income developing countries increases inequality in both groups. Meschi and Vivarelli interpret their finding as evidence for the introduction of SBTC from developed to developing countries. However, no measure is included of the technologies transferred or the transmission channels through which wages are affected, a concern which has also been raised by Conte and Vivarelli. The present paper therefore differentiates trade flows by technology, thereby measuring the inherent skill content of trade. This enables the testing of whether it is indeed more skill-intensive technology which increases wage inequality through technology transfer.

4.2.3.3 Summary and Hypotheses

The main innovation of this paper vis--vis the existing studies is the introduction of an index of technological change, representing the most important control variable in the empirical analysis. The paper furthermore expands on existing studies in three ways: i.) it employs a comparative advantage-based rather than an income-based country classification; ii.) it classifies trade flows according to their technology content - measured by the degree of human capital necessary to produce the goods, which allows for the testing of which types of technology matter most for wage inequality in developing countries, and whether it is indeed more skill-intensive technology which raises inequality the most; and iii.) it uses a consistent version of the Theil index of inter-industry wage inequality, based exclusively on the UNIDO industrial statistics, with comparable values over time containing the same sectors every year (see chapter 3). The sectoral classification used for the computation of the inequality index is the same as for the trade data as well as the classification of industries into different skill levels, and is therefore able to capture the effects of trade on between-sectoral inequality in manufacturing quite well.

A number of hypotheses regarding the effects of different types of traded goods, different groups of trading partners, and different receiving countries can be derived from the literature. For aggregate trade flows, a simplistic view of developing countries would hypothesize an equalizing impact of trade on wages. Technology transfer might have opposing effects, conflating the negative impact and rendering a prediction on the overall impact difficult. Finally, trade could really be driven by domestic technological change, and hence the effect might diminish with the inclusion of the control variable.

As for different levels of technology, once can expect a disequalizing (i.e., positive) impact for trade in higher levels of technology, both due to technology transfer as well as H-O effects. I expect low-technology trade to decrease inequality through exports, while no effect should be present for imports if one believes the predictions of the specific factors model.

The next set of hypotheses pertains to the differential impacts in countries of different relative skill endowments. I use country group interactions to test whether the effects of trade on wage inequality differ between particular groups of countries. South-South trade theory suggests that medium-low technology exports have a particularly strong disequalizing impact in the medium-education countries, since this is where their comparative advantage is located. The "absorptive capacity" argument would furthermore suggest that technology transfer effects of medium- and high-technology imports are stronger (in absolute terms) in the more educated trading partners, i.e., UMECs and LMECs. Overall, I therefore expect to find strong disequalizing effects of medium-low technology in the LMEC and UMEC groups. A structured summary of the hypotheses derived for the different country groups is presented in table 4.1.

Table 4.1: Hypotheses of effects on inequality by technology level and country group

	H-O theory (SS trade version)	Specific factors model	Technology transfer	Overall
High technology exports	+	+	+ (in UMEC)	+
Medium-low tech. exports	+ (in UMEC/LMEC)	++	+ (in UMEC/LMEC)	++
Low technology exports	-	?	?	-
High technology imports	?	–	+ (in UMEC)	-
Medium-low tech. imports	?	-	+ (in UMEC/LMEC)	?
Low technology imports	?	?	?	?

Notes. HEC=High-education country (highest quartile), UMEC=Upper-middle education country (third quartile), LMEC= Medium-low education country (second quartile).

4.3 Empirical Analysis

4.3.1 Data and descriptive statistics

4.3.1.1 Country classification

As has been derived from the literature on "South-South" trade, it is important to distinguish between different types of developing countries to arrive at clear predictions about the effects of trade on wages. Countries are typically classified into different levels of development according to their income, as in the widely used World Bank classification based on GNI. In the context of this analysis, a classification by relative endowments - i.e., the skill-level of the labor force - is more appropriate. Relative human capital endowments are the source of comparative advantage in trade and hence the relevant characteristic from which to derive hypotheses about the impact of trade on wage inequality. Studies supporting this approach are Gourdon et al. (2008), who test H-O theory by introducing interactions with country endowments and find supporting evidence for its predictions, and Forbes (2001), who directly tests different country classifications. She concludes that any classification based on comparative advantage (years of education, wages, or a mix of the two) performs superior to income-based classifications in that the presumed effects of trade are found with the former classification, whereas the latter one yields only insignificant coefficients.

Human capital is proxied by average years of schooling of the population aged 25 years and older, extracted from Barro and Lee (2013a) and extrapolated for the years missing between the 5-year intervals in which the original data are reported.[3] As it is relative endowments that should matter for trade, countries are grouped into quartiles. In previous analyses, developing countries were divided into two or three groups of low-, lower-middle and/or upper-middle income countries according to their per capita incomes, following the World Bank classification. Translating these groups into education, the resulting classification divides countries into low (LEC), lower-middle (LMEC), upper-middle (UMEC), and high (HEC) education. The lower 3 quartiles are considered "developing" and form the estimation sample. Countries classified as HEC are used for classifying trade flows in order to capture technology transfer from more developed countries, and then removed from the sample. Of the 60 countries and total of 1151 country-year observations used in the estimation sample, 20 percent are classified as LEC, 41 percent as LMEC and 39 percent as UMEC. For every developing country, all trade flows to and from countries classified as HEC are summed up. The same is done for the other income categories, so that the South-South hypothesis of trade between developing countries can be tested. The disaggregated trade variables are denoted by affixes numbered 1 to 4 according the trading partner's relative education level from low to high education respectively. They

[3]As noted by Wößmann (2000), years of schooling are not a ideal measure for skills without taking quality of schooling into account, which not only varies greatly between countries, but also over time. It is even more contentious to equate formal schooling with human capital, which has many other components besides education. However, alternative measures for human capital are scarce and those for schooling are equally contested. While there have been attempts to measure educational outcomes directly via cognitive tests (e.g. the "Schooling Quality in a Cross-Section of Countries" dataset by Barro and Lee (2013b), the resulting data are sparse and would virtually eliminate the present panel.

are further decomposed into their technology content as explained in the following section.

4.3.1.2 Trade and technology

The data on trade consist of the total value (in billions of US dollars) of yearly bilateral trade flows between country pairs, provided by the UN Comtrade database.[4] Traded products are coded according to their technology level. The technology classification is taken from Loschky (2010), who calculates R&D intensities of product groups at the ISIC Rev. 3 level.[5] Three categories of technology intensity are employed: Low technology (LT), medium-low technology (MT), and medium-high to high technology (HT). Aggregation is again carried out by adding up the total value of yearly trade in each technology category, separately for imports and exports.

The following graphs depict some basic trends in the trade data. Figure 4.1 shows the rise in developing country trade (estimation sample average) in billions of USD over the sample period. Trade has grown tremendously between 1970 and its peak value in 2010. The share of trade with relatively more high-technology intensive goods has also risen over time, as is apparent from figure 4.2.

[4]Because the trade data are not available in the ISIC scheme, they have to be converted from the Standard International Trade Classification (SITC) using correspondence tables. While a direct conversion is possible for post-1987 data which is provided in the SITC Rev.3, data from 1970 are only available in ISIC Rev.1, for which there is no direct correspondence table to ISIC Rev.3. The data therefore have to first be converted into the SITC Rev.3, and then further into the ISIC classification. Correspondence tables are taken from the EU RAMON database. Conversion is always based on the most detailed (5 digit) product level, whereas the trade data is provided at all levels of aggregation. However, "The values of the reported detailed commodity data do not necessarily sum up to the total trade value for a given country dataset. Due to confidentiality, countries may not report some of its detailed trade. This trade will, however, be included at the higher commodity level and in the total trade value." (Comtrade 2014). After conversion, whenever a higher commodity level trade value deviates from the sum of its sublevel trade value and the higher level contains different sub-level technology groups as per the official classification scheme, a precise recording and grouping of all data is not possible. Hence, only data provided at the 5-digit level is retained so that all the data can be coded into technology levels.

[5]Although Loschky (2010) differentiates between low-, medium low-, medium high-, and high-technology, the upper two categories are pooled together. This is done for two reasons: (1) Retaining consistency with the classification of industries used in the dependent variable, which is based on the 2-digit level of ISIC Rev. 3. The distinction between medium-high technology and high technology is made on a deeper level of product classification which often involves four digits, and pooling the top categories together avoids the resulting overlaps of medium-high and high technology sectors in the wage inequality measure. (2) The trade share of the combined category is already relatively small (around 20% on average), so separating between the categories would lead to more missings, thereby aggravating country composition effects and further complicating the analysis with the introduction of a fourth category.

Figure 4.1: Total trade (country average), in constant (2005) bn. USD

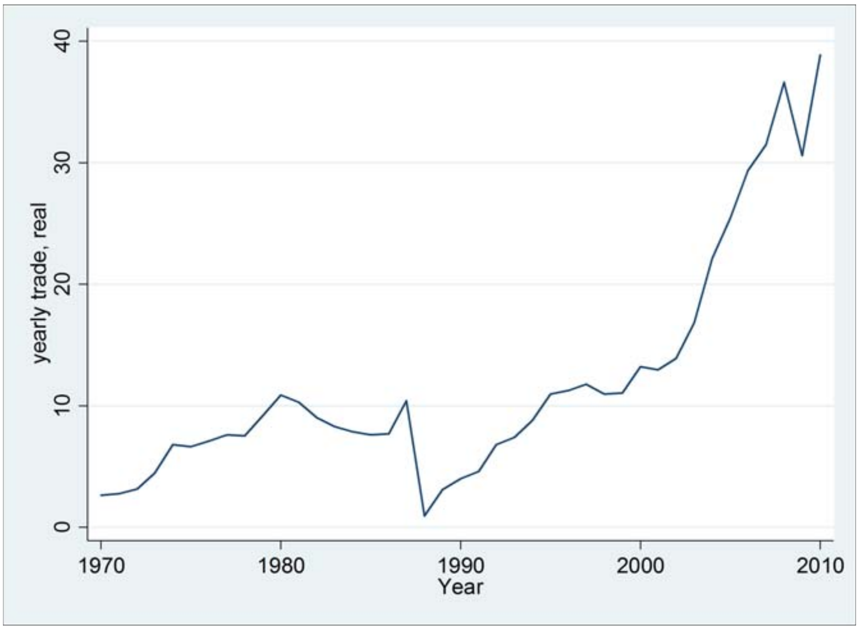

Notes. Nominal USD values from comtrade have been deflated using the US GDP deflator from the WDI. Total trade is the sum of exports and imports.

Figure 4.2: Trade by technology levels (imports and exports)

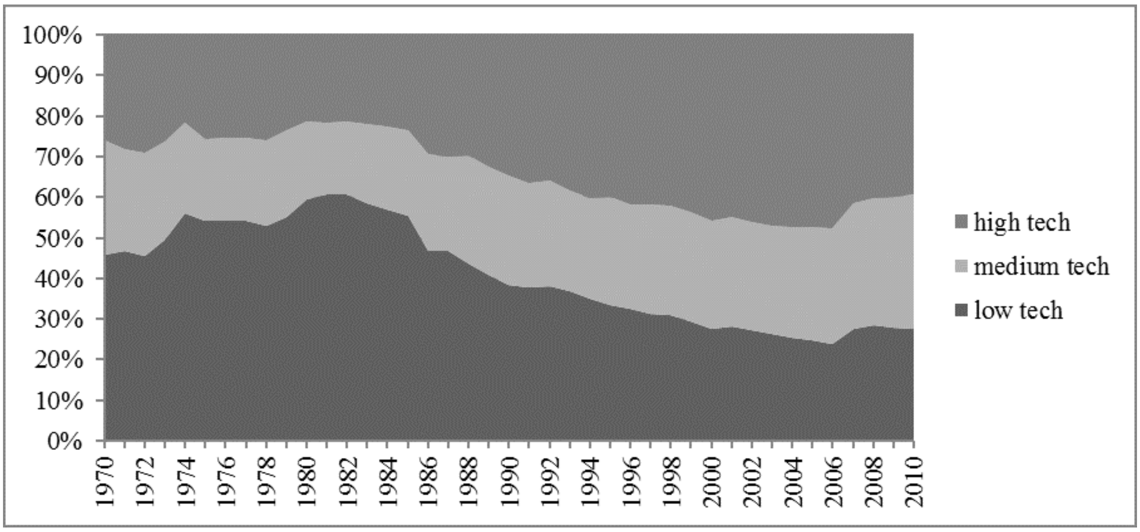

4.3.1.3 Inequality: a sectoral approach

This paper considers the effects of trade on wage inequality rather than income inequality, which is more frequently analyzed in the literature. This more narrow focus has several advantages for the purpose of this paper. It is closer to the theoretical argument that the influence of trade and technology on inequality works via their impact on skill premia. Skill premia directly affect the wage structure, but presumably have a weaker impact on overall income, which has many more components besides wage income and where household formation and composition plays an important role. One would have to identify the impact of trade on the return to other production factors such as capital and land which are both a source of comparative advantage in international trade and a component of income. Also, wage data are more comparable across countries than the available income data, which differ considerably in both quality and content both between countries and over time.

A Theil index of between-sectoral wage inequality has been constructed to serve as the dependent variable in the empirical analysis. The index is based on the UNIDO industrial statistics on manufacturing, using data from 1970 to 2010. Although a similar index has been built by the University of Texas Inequality Project (UTIP), it is not clear which data enter their index, as the raw data require several choices as to which sectors to include in order to retain consistency and ensure comparability over time. Hence, the index has been re-calculated for the entire time period. Different versions of the index are employed to test the robustness of the results to the choices made in obtaining a consistent inequality measure. A discussion of the advantages and weaknesses of the sectoral approach using the UNIDO data vis--vis Deininger and Squire's (1996) more frequently used individual-based dataset of Gini coefficients can be found in Conceição and Galbraith (2000). The main results are robust to using the UTIP index rather than the newly calculated index, as discussed in the robustness checks.

Similar to the technology classification, the UNIDO statistics are also based on the ISIC sectoral classification and thus match the trade data perfectly. The entire analytical set-up is based on a sectoral approach. It hence captures sector-biased ("asymmetric") rather than "simple" factor-biased technological change which affects all sectors of the economy to more or less the same extent (symmetric). There are two reasons for choosing the sector-based approach. First, the technology content of trade flows is measured by the technology content of the traded goods, which is based on the classification of the respective industry from low- to high technology. This measure does not capture differences in the within-industry composition of skills - it can therefore only explain changes in the distribution of wages between industries, which is what the inequality index measures. Second, a sector bias of skills is a much more reasonable assumption than simple factor bias, especially if one drops the unrealistic assumption of the homogeneity of labor. A highly qualified worker in the metal working industry is most likely to have different kinds of skills than a highly qualified worker in, say, the apparel industry. Even though they may have the same level of qualification, the wage premia of the two are likely to be driven up to a different extent by factor-biased SBTC. Similar to the terminology used by Haskel and Slaughter

(2002), the term sector-biased SBTC is used here to include not only the obvious sector-specific SBTC, but also the pervasive, asymmetric factor-biased SBTC because it affects some sectors more than others.[6]

One drawback of the sector-focused approach is that factor-biased SBTC which affects sectors asymmetrically can be conflated in the computation of industry wage averages, which the employed between-sector inequality measure relies on. The problem arises because the skill-composition of the workforce varies between sectors. A numerical example for the problem can be found in part B of appendix 4. However, there is little reason to suspect that results will be distorted systematically, and the between-unit measure can be interpreted as the lower bound to overall inequality (Conceição and Ferreira 2000)

The dataset resulting from the construction of the Theil index contains more than 3000 observations over the years 1970-2010, but the observations and countries covered are reduced substantially in the course of the sample construction. The between-sector component of the Theil is defined as

$$T' = \sum_{s=1}^{S} y_s \cdot ln(\frac{y_s}{n_s})$$

with S denoting the different sectors, s=1, ..., S. ys represents the wage share of sector s, defined as the sector average over the total average wage of all industries. ns represents each sector's population share, defined as the sector's population Ns over total population N (cf. Theil (1967): 95). The measure is not uniformly bounded upwards, which makes intuitive interpretation of its values difficult. It therefore enters the regression in log-specification to ease interpretation. The development of the (in-sample) Theil index over the sample period (1970-2010) is displayed in figure 4.3.[7] As with trade, there is a clearly discernible upward trend over time.

4.3.1.4 Control variables

Technological change

The difficulty with including technological change in empirical analyses is measurement. Even though efforts have been made to find appropriate proxies, technological change is often simply defined as the unexplained residual of wage determination models. As argued by Topel (1997: 60), this "makes it nearly impossible for [the theory that technological

[6] While there are several theoretical analyses on the effects of factor- vs. sector-biased SBTC on wages (see, e.g., the studies referred to by Haskel and Slaughter 2002. Stehrer (2010) points out that the results depend on the specific assumptions of the theoretical models and there is no conclusive overall result. Unfortunately, there are only few studies that empirically examine the importance of sector- vs. factor-biased technical change and they are limited to developed countries. The results do, however, all indicate an important role of sector-biased SBTC in explaining relative wages. Haskel and Slaughter (2002) conclude that the sector bias of SBTC is the decisive factor in explaining changes in skill premia, but they also find a smaller role for a factor bias. De Santis (2002) also finds in his analysis of a general equilibrium model with H-O-trade applied to US and UK data that sector-biased technical change performs relatively better than factor-biased technical change in explaining the data.

[7] Note that in line with the explanations in chapter 3, the "dynamic" version of the index developed there is used.

Figure 4.3: Development of the Theil index of inter-industry wage inequality

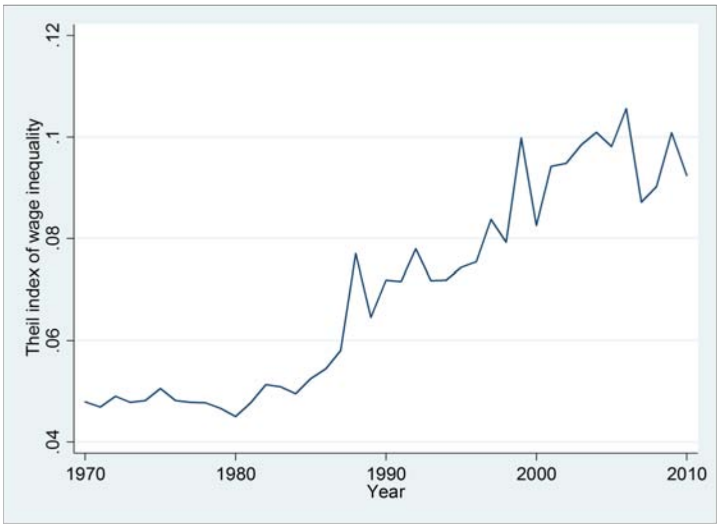

Notes. The graph is based on the estimation sample of countries.

change, altering the demand for the two kinds of labor by changing their relative productivities, is responsible for an increase in wage inequality] to fail." An attempt to find a measure of technological change has been made by Jaumotte et al. (2013), who use the share of domestically produced information and communications technology capital in the total capital stock. The variable turns out to significantly increase inequality in both developed and developing countries while trade itself has an equalizing effect on the income distribution. However, technological change in developing countries is likely to start at much less sophisticated levels of technology, which this measure does not capture. Technological change would consequently be underestimated. Zhu and Trefler (2005) use labor productivity to measure technological change and also find a positive relationship with trade. Gourdon (2011) argues that total factor productivity (TFP) would be more appropriate but also uses labor productivity in his analysis because of better data availability. Lipsey and Carlaw (2004) challenge the interpretation of TFP as measuring technological change. They argue that positive changes in TFP simply reflect the surplus returns that emerge from investing in new technologies which are necessary to recoup the investment. Consequently, if there are no surplus returns, technological change goes unmeasured. Nevertheless, although it may underestimate the true extent of technological change, TFP-based measures are the best feasible option given the data available. As long as the unmeasured components of TFP are not occurring systematically, this merely adds noise to the data.

To arrive at a measure of technological change, I calculate a productivity index which decomposes observed changes in the input-output ratio of production into different components. Besides the different aspects of technical and scale efficiency, this also entails a component of technical change, capturing movements in the production frontier. Data Envelopment Analysis (DEA) is employed to estimate the technological frontier, defined as the maximum level of TFP observed in all the production units of the data. The DPIN program (V.3), developed and provided by O'Donnell (2014), uses linear programs for estimation. Of the several available productivity indices, a Färe-Primont index is chosen since it fulfills the transitivity criterion by which obtained values can be meaningfully compared across time as well as production units.[8] The UNIDO data, which have partly already been used in the inequality index, are exploited again for the calculation of the index. Besides wages, the dataset also contains information on capital, output, and value added. In order to not get biased results due to unaccounted intermediate inputs, value added rather than output is used as the output measure, and both wages and capital are included as inputs. Unfortunately, the data on capital are scarce, and using the TFP technological frontier reduces the sample by 40%, despite the imputation of missings as described below. The index is therefore estimated again measuring only labor productivity. The same procedure as for the TFP index is applied, but using only labor as an input. The correlation analysis between the total- and the labor-productivity indices for those cases where both are available suggests that they capture the same movements of the production frontier in all but a few countries. Hence, the labor productivity index is used in the preferred specifications as it results in wider country coverage, and the TFP index is employed as a robustness check, yielding similar results. As the data are reported at the sectoral level, sectors are "production units" in the estimation of productivity.[9] The technically most efficient sector determines the production frontier, which is then used as the control variable for technological change in the regressions. Three different versions of the index are constructed, which use different sectors and imputation methods for missing values: One wherein missing sectors are substituted for by other sectors (imputation across sectors), one wherein the same procedure is applied but only those sectors which have less than 50% missings are used, and one wherein all sectors are used and missings are substituted for with values from the same sector in earlier years.[10] The index relying on cross-imputed values is used in the preferred estimations as it adds no "new" information from other years to the data in a given year. As a robustness check, the other two indices are tested as well and the results show that they yield virtually the same estimates (see table 4.7).

Labor supply

Value added in agriculture is included as a supply-side control variable in the spirit of Lewis' (1954) dual-sector model. The variable is supposed to measure the amount of unskilled surplus labor in an economy, which might prevent wages at the very bottom of the distribution from rising despite increased demand through trade and/or technology. The data come from the World Bank's World Development Indicators (WDI). Value added in agriculture is chosen over the share of employment in agriculture, which seems closer to the labor supply it is supposed to capture, and has been used by, e.g., Jaumotte et al.

[8]The Färe-Primont index has been developed by O'Donnell (2014) and is based on the ratio of two versions of the (more commonly known) Malmquist index developed by Färe and Primont (2012).

[9]Productivity is estimated separately across country, as the DPIN program does not allow a multi-level equation system (country and sectoral level). Values can therefore only be meaningfully compared within a country over time. Though the within-estimator is used in the empirical analysis, this does not represent a problem here.

[10]Values from earlier years are used in order to not overestimate technological progress, which can reasonably be assumed to evolve positively over time. Values from subsequent years are only used in the exceptional cases where no values are available for previous years.

(2013), due to a greater country coverage. In preliminary tests on the data, the two measures produce the same results.

Human capital

Although countries have already been grouped according to their relative human capital endowments, education levels still matter as they constitute a (short-term) measure of the supply of skilled labor, which can mitigate pressure on high-skilled wages, and reduce skill premia. The same linearly interpolated Barro and Lee (2013a) data are used as for the country classification.[11]

FDI

Inward FDI flows (taken from UNCTAD) are included (in Mio. USD) in order to control for an alternative source of technology transfer likely to be correlated with trade. The direction and form of the effect has not been established unambiguously in the literature (on a review of recent results from empirical studies, see Figini and Gorg 2011). However, since the assumption that FDI influences inequality via skill premia follows the same line of argument as the hypotheses on the effects of trade, the variable has been frequently included in analyses on the effects of trade on income inequality (e.g., Jaumotte et al. 2013, Gourdon 2011, and a number of country case studies) and has been often found to significantly increase inequality.

GDP

GDP is included in order to control for "size-effects": All other things equal, richer economies trade more in absolute terms and hence without taking economic size into account, one might hypothesize that larger countries are always more (un-)equal, depending on the assumed effect of trade on inequality. On the other hand, larger economies tend to trade less in relative terms due to a larger domestic market, so the overall effect remains unclear. Real expenditure-based GDP in (2005) PPP adjusted USD is taken from the Penn World Tables, Version 8.0 (Feenstra et al. 2015), and the variable enters in logarithms.

A list of the countries in the sample, as well as the in-sample means of the most important variables can be found in appendix table 4.A.1.

4.3.2 Model specification

The basic model has the following functional form:

$$ln(Theil_{i,t}) = \alpha + \rho ln(Theil_{i,t-1}) + \beta Trade_{i,t-1} + \lambda Technology_{i,t-2} + \sum_k \gamma_k X_{k,i,t} + \delta_i + y_t + \epsilon_{i,t}$$

Indices t and i denote year and country, respectively. Trade covers the different specifications of the trade variable (e.g., interactions with country dummies, separate consideration of imports and exports), which enters the model with a one-period lag to allow for a time lag in the adoption of imported technology.[12] X is the set of k control variables, all of

[11] The fact that the same measure is used does not affect the estimates, neither for the aggregated, nor the disaggregated (education-based) country group data. In fact, the impact of the education variable on the coefficients of interest is negligible. Results are available upon request.

[12] The inclusion of the trade variable with a lag of 1 period is chosen for several reasons. Descriptive correlations between trade in different technology levels and the inequality measure suggest that the first

which enter the regression in levels, and some of which are lagged by one (FDI) or two (technological change) periods. Both country fixed effects (δ_i) and time dummies (y_t) are included. $\epsilon_{i,t}$ denotes the usual error term.

Even though the inter-industry Theil index exhibits less inertia than other measures of income inequality such as the Gini index, misspecification tests in a static model indicate the presence of autocorrelation. A dynamic specification is therefore appropriate. The dynamic fixed effects OLS model delivers biased estimates (primarily of the lagged dependent variable) in a finite sample due to the correlation between the lagged dependent variable and the error term as described by Nickell (1981) and therefore referred to as "Nickell bias," or LSDV bias. Although alternative (IV-based) estimation techniques are available for dynamic panel models, the most widely used being the Generalized Method of Moments (GMM) (Arellano and Bond 1991), the preferred specification here is the simple FE model. Tentative faith is put in these estimates for two reasons: First, the LSDV bias is a problem of small T, and although an average of 16-19 years is not yet "large T" (starting from around 20 years), it is not considered small either. Second, while the bias is quite severe in the autoregressive (AR) term, it is much smaller for the ""-variables, i.e., all other ("control-") variables in the model. Results from several simulation studies suggest that the bias amounts to less than one percent of the coefficient estimate given the values of and T in the panel at hand (e.g., Judson and Owen 1999; Köhler et al. 2011). A robustness check using GMM is nevertheless conducted, indicating that the LSDV bias is not a problem in the present sample given that the more precisely estimated coefficients change very little between fixed effects and the GMM specification and even increase in several instances.

4.4 Results and Discussion

For testing hypotheses about the impact of trade on wage inequality in different country groups, at different technology levels, and from different trading partners, many possible specifications can be employed. At the most disaggregated level of the trade data and with the introduction of the country dummies, the number of trade variables would rise to 72, which is not operational given that the number of cross-sections is around 60. The approach taken is to start from the most aggregated level and to move stepwise to more disaggregated specifications. Total trade values are investigated first, before moving to exports and imports separately. Each group is further disaggregated by technology, and differential impacts in countries of different relative education levels are tested in the next step.

The technological change variable is included with a two-year time lag in the preferred specification. This is done because in its contemporary version, technological change is likely to be influenced by trade itself, which also enters the model with a one period lag. The impact of the variable is interpreted as follows: If the coefficients are affected by the inclusion of the variable, this means that the observed effects on wage inequality

lag is the most relevant one. Furthermore, most of the literature has used one-period lagged trade variables. Lastly, the inclusion of further lags would significantly reduce the estimation sample.

are possibly not due to trade, but rather that both variables are at least partly driven by domestic technological change. The effect can of course also go the other way, i.e., technological change can be disequalizing and generate trade flows which have per se an equalizing impact, in which case the two opposing effects may become apparent only after technological change is controlled for. H-O theory does not yield any predictions about the effect of imports on the distribution of factor rewards - they are merely the mirror image of a country's specialization according to its comparative advantage, which is reflected in the export structure. The specific factors model on the other hand would suggest that imports may have distributional effects. In particular, skill-intensive imports may lower the gains from trade for those at the upper end of the wage distribution and thereby mitigate disequalizing technology transfer effects which exert pressure on skill premia.[13]

4.4.1 Aggregate Trade

Table 4.2 shows the results for the most basic specification, where all trade flows (imports and exports) have been added up.[14] Trade is significantly negatively related to wage inequality, but in terms of economic size, the effect is rather small. According to the coefficient estimate, a 1 billion dollar increase in trade would reduce wage inequality by little over 0.04 percent. The effect persists with the inclusion of the control variable for technological change (column 2), but is reduced substantially and becomes insignificant. This indicates that a major part of the effect on wage inequality attributed to trade might indeed stem from technological change, which, in line with expectations, leads to higher wage inequality and is significant at the five percent level. As for the remaining insignificant control variables, FDI has the expected positive coefficient, but the sign on the education variable is not in line with expectations: a higher (short-run) supply of skilled labor, as captured by the "years of education" control variable, does not seem to lower the pressure on skill premia. Rather, even within the already more homogeneous country groups, a better educated workforce is associated with more inequality. Although insignificant, the fact that the variable is reduced substantially by the inclusion of the technological change control variable suggests that this effect could have something to do with its absorptive capacity, wherein a more educated workforce is more apt to adopt (inequality-increasing) technology. A higher share of unskilled workers, proxied by the share of value added from agriculture in GDP, does not seem to have any appreciable impact on wage inequality given the small and volatile coefficient. Finally, a higher GDP seems to be associated with lower inequality, which might have to do with the fact that large economies trade relatively less, but the effect is not significant, either.

To check whether the effect is driven by trade of a particular technology intensity, trade flows are decomposed into low-, medium-, and high technology in columns 3 and 4 of table 4.2. Although none of the coefficients are significant, it is interesting to again note the substantial decrease in the coefficients on all three variables once the technological change

[13]Low-skill intensive imports on the other hand could theoretically increase wage inequality since they compete with local industries employing labor from the lower end of the wage distribution. However, low-skilled labor is arguably much less specific than skilled labor and hence less susceptible to such effects.

[14]Including imports and exports separately does not yield any new insights, with neither variable being individually significant. Results are available upon request.

Table 4.2: Results total trade

Dep. var.: ln(Theil)	(1)	(2)	(3)	(4)
ln(Theil)(t-1)	0.789***	0.781***	0.788***	0.781***
	(0.0335)	(0.0365)	(0.0342)	(0.0368)
Tech(t-2)		0.244**		0.243**
		(0.0915)		(0.0917)
Totaltrade(t-1)	-0.000415*	-0.000105		
	(0.000237)	(0.000509)		
Lowtech(t-1)			-0.00152	0.000747
			(0.00229)	(0.00328)
Medtech(t-1)			0.00107	-0.000814
			(0.00193)	(0.00320)
Hightech(t-1)			-0.000722	-0.0000161
			(0.000563)	(0.000889)
GDP	-0.0662	-0.0280	-0.0692	-0.0258
	(0.0716)	(0.0865)	(0.0759)	(0.0910)
Edcuation	0.0189	0.00790	0.0199	0.00763
	(0.0352)	(0.0347)	(0.0351)	(0.0353)
ValAddAgri	0.00224	-0.000368	0.00213	-0.000188
	(0.00414)	(0.00429)	(0.00427)	(0.00463)
FDI(t-1)	0.00290	0.00290	0.00252	0.00299
	(0.00414)	(0.00322)	(0.00415)	(0.00338)
Quartile2(LMEC)	0.0919	0.0915	0.0926	0.0903
	(0.0678)	(0.0852)	(0.0681)	(0.0858)
Quartile3(UMEC)	0.119	0.115	0.120	0.114
	(0.0819)	(0.1000)	(0.0825)	(0.0999)
Observations	1,151	903	1,151	903
R^2	0.689	0.679	0.689	0.679
# of countries	60	58	60	58
Year FE	YES	YES	YES	YES

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1. LMEC = lower middle education country, UMEC = upper middle education country.

control is included in column 4. The sign on the low- and medium-technology coefficients even reverses. This again provides indication that the previously discussed problem of omitted variable bias is present, and that controlling for technological change is important in order not to falsely attribute technology effects to trade.[15]

4.4.2 Disaggregated Results: Technology

Next, imports and exports are considered separately, while at the same time retaining the three different technology levels. Columns 1 and 2 of table 4.3 show the export regressions and columns 3 and 4 the import ones. Total imports (exports) are included as a control variable, and the first and second column of each panel contain the estimates with and

[15] For simplicity reasons, the coefficients will not be shown in the remaining tables. Instead, the top row will indicate whether the technological change variable is included in the model.

without the technological change control, respectively, as also indicated in the top row. The full set of the previously discussed control variables is included, but omitted from the table for simplicity.[16]

Table 4.3: Results imports and exports by technology levels

Dep. Var.: ln(Theil)	Exports		Imports	
	No tech	Tech(-2)	No tech	Tech(-2)
Lt_trade(t-1)	-0.00462	-0.00119	0.00827	0.00206
	(0.00306)	(0.00352)	(0.00801)	(0.00864)
Mt_trade(t-1)	**0.00692****	0.00337	-0.00968	-0.00512
	(0.00313)	(0.00448)	(0.00600)	(0.00583)
Ht_trade(t-1)	-0.000422	-1.11e-05	-0.00122	-0.000582
	(0.00143)	(0.00147)	(0.00148)	(0.00199)
Totalimp/exp(t-1)	-0.00187	-0.00158	0.000544	0.000712
	(0.00132)	(0.00147)	(0.000853)	(0.000498)
Observations	1,151	903	1,151	903
R^2	0.690	0.679	0.690	0.679
Number of countries	60	58	60	58
Control variables	YES	YES	YES	YES
Year FE	YES	YES	YES	YES

Notes. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Ht = high technology, mt = medium-low technology, lt = low technology.

A few interesting results emerge. First, the coefficient on high-technology trade remains negative, small, and insignificant for both exports and imports. Second, the signs on both medium- and low-technology trade flows are opposite for imports and exports. The negative coefficient of low-technology exports as well as the positive sign on medium technology exports are in line with expectations and the hypotheses derived from South-South trade theory, but remain insignificant. Third, while it seems that medium-low technology exports are significantly associated with increasing wage inequality, this effect disappears with the inclusion of technological change in column 2. It would appear that technological change is driving at least part of the positive coefficient of medium-low technology trade - in fact, the coefficient estimates suggest that around half of the effect is due to technological change rather than trade. Finally, it is also worth noting that once the trade flows are disaggregated a bit more, some of the coefficients are substantially larger than those on the aggregate trade variable in table 4.2. For example, the coefficient estimate of column 1 would imply that a billion dollar increase in medium-technology intensive export is associated with approximately 0.7 percent lower wage inequality (although, as just discussed, the effect actually attributable to trade is only around half of that).

[16] The estimates for the control variables for columns 2 and 4 can also be found in appendix table 4.A.4, columns 1 and 3.

4.4.3 Disaggregated Results: Country Groups

As summarized by the hypotheses in table 4.1, the next set of regressions uses country interactions to test whether differential effects materialize in particular groups of countries. The lack of results for high-technology trade, for example, could be due to opposing effects in different country groups which offset each other, which can be disentangled in this more differentiated set-up. I then test whether these can be attributed to imports and exports separately. Estimation results are presented in table 4.4 and are ordered according to technology, starting with high technology trade. Again, the control variables are included, but not shown. Results for total trade, comprising both imports and exports, is shown in the first panel, and exports and imports are separately accounted for in panels 2 and 3, respectively.

In line with the results from the previous set of regressions, the effects of high-technology exports are not only nowhere near significance, but also smaller than the coefficients on other technology groups throughout all regressions. In particular, contrary to what has been found in the previous literature, there is no evidence for a disequalizing technology transfer through imports in the more educated country groups, where the coefficients are in fact negative.

The results from columns 1 and 2 indicate that the previously found positive coefficient on medium-technology trade mainly occurs in the more educated country groups, which is in line with the South-South trade logic and the adaptive capacity argument. Both exports and imports have negative signs, but are not separately significant. Despite not being significant, the results on medium-low tech exports confirm the previous patterns: the coefficients are highly affected by the inclusion of the technological change control variable, with the coefficient increasing substantially for LMECs, and turning positive for UMECs. No such effects are found for imports, rendering credibility to the story of a third variable bias through technological change, which works via exports.

For low technology trade, surprisingly, a significant disequalizing association emerges in the least educated country groups. When consulting columns 3-6, it is clear that this effect arises mainly through low-technology imports, where the coefficients are larger and the positive coefficient in LECs is significant as long as technological change is not included. One explanation for the negative sign could be import competition; another one could be the introduction of labor-saving technology. The lack of results pertaining to low technology imports once technological change is controlled for supports the hypotheses derived from the H-O and the specific factors models, as they do not predict any effects for low-skill (and hence unspecific) factor intensive imports. The negative coefficients in the other country groups on the other hand are in line with H-O theory, with clearer results for exports, which are significant for LMECs.

Another point worth mentioning is that again, coefficients increase compared to the previous, more aggregate specification. It seems that the more detailed the specification, the more it is able to capture the various heterogeneous effects of trade flows, which, at the aggregate level, cancel each other out and lead to a very small overall coefficient. To provide an example of the size of the effect, the significant coefficient in lower-middle

Table 4.4: Results by technology level and country group

Dep.var.: ln(Theil)	(1) Total trade		(2) Exports		(3) Imports	
	No tech	Tech(-2)	No tech	Tech(-2)	No tech	Tech(-2)
Ht_trade(t-1)	0.0107	0.0121	-0.0232	-0.0132	0.0137	0.0166
	(0.0175)	(0.0154)	(0.0287)	(0.0296)	(0.0195)	(0.0203)
LMEC*ht_trade(t-1)	-0.0121	-0.0140	0.0205	0.00279	-0.0147	-0.0165
	(0.0176)	(0.0154)	(0.0299)	(0.0318)	(0.0195)	(0.0200)
UMEC *ht_trade(t-1)	-0.0104	-0.0120	0.0244	0.0127	-0.0137	-0.0166
	(0.0177)	(0.0158)	(0.0292)	(0.0300)	(0.0196)	(0.0205)
Mt_trade(t-1)	-0.0181	-0.0196	0.00300	-0.0124	-0.0371	-0.0338
	(0.0129)	(0.0128)	(0.0139)	(0.0178)	(0.0232)	(0.0252)
LMEC*mt_trade(t-1)	0.0198	0.0263*	0.0105	0.0401	0.0227	0.0299
	(0.0130)	(0.0137)	(0.0183)	(0.0259)	(0.0244)	(0.0269)
UMEC*mt_trade(t-1)	0.0174	0.0170	-0.00361	0.0101	0.0301	0.0253
	(0.0132)	(0.0130)	(0.0150)	(0.0178)	(0.0247)	(0.0269)
Lt_trade(t-1)	0.0185*	0.0178*	0.0113	0.0181	0.0468*	0.0359
	(0.00961)	(0.00973)	(0.00957)	(0.0125)	(0.0273)	(0.0304)
LMEC*lt_trade(t-1)	-0.0193**	-0.0229**	-0.0195*	-0.0310**	-0.0363	-0.0387
	(0.00944)	(0.0109)	(0.0109)	(0.0153)	(0.0278)	(0.0328)
UMEC*lt_trade(t-1)	-0.0194*	-0.0163	-0.00949	-0.0146	-0.0501	-0.0383
	(0.0100)	(0.0111)	(0.0107)	(0.0147)	(0.0305)	(0.0353)
Observations	1,151	903	1,151	928	1,151	903
R^2	0.689	0.680	0.690	0.302	0.691	0.680
# of countries	60	58	60	58	60	58
Control variables	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES

Notes. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Ht = high technology, mt = medium-low technology, lt = low technology. LMEC = lower middle education country, UMEC = upper middle education country.

education countries indicates that a billion dollar increase in low-technology exports is associated with a decrease in wage inequality of around 3 percent. Given the yearly mean of 5.7 billion for LMECs, this is a potentially rather powerful equalizer. It is also worth noting that again, the import coefficients are affected much less than the export ones by the inclusion of the technological change control variable, supporting the view that the variable indeed captures what it is supposed to: domestic technological change, rather than technological change through trade.

Summary of results

Summing up the insights obtained from the regressions, one can extract four main findings from the many results. First, although the coefficients are mostly insignificant, low-technology trade seems to be generally equalizing, as predicted by H-O theory.

Second, medium-low technology exports have positive coefficients and seem to be disequalizing in all but the countries in the lowest education quartile. This finding fits with the South-South trade story as well as the technology transfer hypotheses, in particular the absorptive capacity argument. Third, there is evidence of technological change driving both exports and inequality in the export regressions, in particular in the medium-low technology sectors and in the more skill-abundant countries, underscoring the need to control for the variable. As expected, the technological change control plays a lesser role for imports, with a lot of coefficients remaining virtually unchanged with the inclusion of the variable, which is in stark contrast to the export results and renders credibility to both the measure and the supposition of omitted variable bias.

Lastly, in contrast to the findings of previous studies, no results emerge for high-technology trade. While it may not be surprising that there are no findings pertaining to exports since little domestic technological advancements in high-tech sectors can be expected in the countries which are relatively less endowed with skilled labor, the fact that there are also no results for imports is surprising. In fact, not only are there no significant positive effects, but the coefficient on high-technology imports is negative throughout all specifications, as well as in upper-middle education countries which arguably are the most apt to introducing such technology. Rather, most technological advancements seems to take place in medium-low technology sectors, both through domestic technological change which also boosts exports, and through technology transfer through imports in the relatively more educated country groups. The latter result should be taken with caution, however, since none of the import coefficients are significant.

4.5 Robustness checks

Although the structure of the present dataset is not ideal for GMM estimation given the comparatively long T of 16-19 years relative to the number of groups (58-60), the method is employed in order to demonstrate that the effect of the LSDV bias on the estimates of the "β"-variables, i.e., the variables of interest, does not change the results substantially. In order to avoid the problem of "too many instruments" (Roodman 2009), the instrument set has been restricted in several ways. The results from difference GMM two-step estimation are shown in columns 2 and 4 of table 4.5, and compared with those obtained using FE in columns 1 and 3. Instruments are restricted to the first few valid lags, and are additionally collapsed in order to keep the number of instruments down. Orthogonal deviations are used in order to mitigate the unbalancedness of the panel. Since the concern here is exclusively with the LSDV bias, only the lagged dependent variable is treated as endogenous.

Table 4.5: GMM results, total trade

Dep. var.: ln(Theil)	(1) FE	(2) GMM	(3) FE	(4) GMM
ln(Theil)(t-1)	0.781***	0.856***	0.781***	0.811***
	(0.0365)	(0.192)	(0.0368)	(0.212)
Totaltrade(t-1)	-0.000105	-0.000141		
	(0.000509)	(0.000714)		
Total_lt(t-1)			0.000747	0.00243
			(0.00328)	(0.00483)
Total_mt(t-1)			-0.000814	-0.00214
			(0.00320)	(0.00476)
Total_ht(t-1)			-1.61e-05	0.000216
			(0.000889)	(0.00174)
GDP	-0.0280	0.00243	-0.0258	-0.0183
	(0.0865)	(0.120)	(0.0910)	(0.155)
Education	0.00790	-0.000308	0.00763	0.000818
	(0.0347)	(0.0358)	(0.0353)	(0.0492)
ValAddAgri	-0.000368	-0.000611	-0.000188	0.00127
	(0.00429)	(0.00509)	(0.00463)	(0.00558)
FDI(t-1)	0.00290	0.00290	0.00299	0.00302
	(0.00322)	(0.00349)	(0.00338)	(0.00344)
Tech(t-2)	0.244**	0.268***	0.243**	0.265***
	(0.0915)	(0.0831)	(0.0917)	(0.0776)
Quartile2(LMEC)	0.0915	0.0806	0.0903	0.149
	(0.085)	(0.101)	(0.0858)	(0.124)
Quartile3(UMEC)	0.115	0.122	0.114	-0.0183
	(0.100)	(0.118)	(0.0999)	(0.155)
Observations	903	845	903	845
R^2	0.679		0.679	
Number of countries	58	58	58	58
Year FE	YES	YES	YES	YES
Number of instruments		56		57
Hansen Test		0.125		0.141
AR(1)		0.00560		0.0115
AR(2)		0.148		0.151

Notes. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Lags have been restricted to lengths 3-10 in column 2, and 5-13 in column 4. The depth of lag lengths has been guided by the misspecification tests. Similar results emerge with varying lag lengths (results available upon request). Ht = high technology, mt = medium-low technology, lt = low technology. LMEC = lower middle education country, UMEC = upper middle education country.

Results show that the negative impact of trade does not vanish when GMM is employed - only some of the coefficients are reduced slightly. Generally, the more precise the coefficient estimate, the more stable it is across different specifications. Some of the more precisely estimated coefficients (most notably, the technological change control variable, but also trade) even slightly increase with the GMM estimator. The coefficient on the lagged dependent variable does increase more substantially, which is in line with the prediction that the LSDV bias entails a relatively larger downward bias on the AR-term. Overall, the results do not provide indication that LSDV bias threatens the validity of the FE estimates.[17]

Table 4.6 contains the estimates obtained when using different versions of the technology index, as described in section 4.3.1.4. Only the results for the preferred specification of the relatively aggregated trade variables are shown here (corresponding to columns 2 and 4 of table 4.2), and the original results using the cross-imputed index are displayed in columns 1 and 4 for comparison. The coefficient of the lagged dependent variable and the country group dummies are omitted. Both the coefficients and the standard errors change very little when the alternative versions of the technology index are used, and the technology indices themselves also yield similar results, although the one using only part of the sectors (columns 3 and 6) is insignificant, which is in line with the fact that it contains fewer sectors and consequently yields less clear results.

Robustness to the TFP technological change index is tested in the following. The estimates in table 4.7 correspond to columns 2 and 4 of table 4.2, which are displayed again here in columns 1 and 4 for comparison. One can see that the result on the aggregate (total) trade variable does not change substantially. A few control variables change signs, most notably the education variable. However, these are very likely to stem from the smaller sample size rather than the difference in the technological change variable, as the results in columns 4 and 5 as well as 7 and 8 suggest, which contain the estimates for each index when executed on a (substantially smaller) common sample. The same can be said about the observed change in the point estimate of the trade variables, which predominantly stem from the difference in sample composition. In fact, the estimates on the small, constant sample yield very similar coefficients on all variables, with the exception of the technological change index itself. It appears that the TFP index is a little more powerful in capturing movements in the technological frontier, as shown by the larger, and more significant, point estimate in the constant sample and the larger effect it has on the trade variables. The true extent of omitted variable bias might therefore be even slightly larger than what was found in the above estimations.

Because results could be more volatile at disaggregated levels, the remaining specifications are checked for robustness as well and displayed in the appendix. Table 4.A.4 contains the results for trade decomposed by the trading partner's relative education classification, and the decomposition by the country group is displayed in table 4.A.5. The original results are displayed in columns 1 and 3 in both tables. Again, the estimates are qualitatively similar between the TFP and the labor productivity index, but are often less significant with the latter. The overall results of testing the TFP-based versus the labor

[17]The GMM results for the remaining specifications can be found in appendix tables 4.A.3 and 4.A.4.

Table 4.6: Robustness of FE results to different labor productivity indices

	(1)	(2)	(3)	(4)	(5)	(6)
Type of imputation	Cross-sectoral	Within-sectoral	Cross-sectoral, only sectors w\>50% data	Cross-sectoral	Within-sectoral	Cross-sectoral, only sectors w\>50% data
Totaltrade(t-1)	-0.000105	-0.000107	-0.000114			
	(0.000509)	(0.000509)	(0.000507)			
Total_lt(t-1)				0.000747	0.000891	0.000845
				(0.00328)	(0.00327)	(0.00328)
Total_mt(t-1)				-0.000814	-0.000919	-0.00103
				(0.00320)	(0.00316)	(0.00318)
Total_ht(t-1)				-1.61e-05	-1.46e-05	5.36e-05
				(0.000889)	(0.000864)	(0.000874)
GDP	-0.0280	-0.0227	-0.0265	-0.0258	-0.0200	-0.0244
	(0.0865)	(0.0869)	(0.0871)	(0.0910)	(0.0915)	(0.0916)
Education	0.00790	0.00931	0.0100	0.00763	0.00903	0.00947
	(0.0347)	(0.0342)	(0.0348)	(0.0353)	(0.0349)	(0.0355)
ValAddAgri	-0.000368	0.000105	-0.000139	-0.000188	0.000320	3.82e-05
	(0.00429)	(0.00435)	(0.00436)	(0.00463)	(0.00471)	(0.00472)
FDI(t-1)	0.00290	0.00293	0.00311	0.00299	0.00302	0.00326
	(0.00322)	(0.00320)	(0.00323)	(0.00338)	(0.00337)	(0.00341)
Tech(t-2)	0.244**	0.231**	0.266**	0.243**	0.232**	0.267**
	(0.0915)	(0.0997)	(0.131)	(0.0917)	(0.0998)	(0.132)
Observations	903	903	903	903	903	903
R^2	0.679	0.679	0.678	0.679	0.679	0.678
# of countries	58	58	58	58	58	58
Year FE & controls	YES	YES	YES	YES	YES	YES

Notes. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Lags have been restricted to lengths 3-10 in column 2, and 5-13 in column 4. The depth of lag lengths has been guided by the misspecification tests. Similar results emerge with varying lag lengths (results available upon request). Ht = high technology, mt = medium-low technology, lt = low technology. LMEC = lower middle education country, UMEC = upper middle education country.

productivity index indicate that for both exports and imports, the coefficients are similar, with occasional changes in significance as well as magnitude, and very few (insignificant) sign changes, of which at least a fraction can be attributed to the smaller sample.

As another robustness check, the Theil index provided by UTIP is used as the dependent variable. Estimation results can be found in appendix tables 4.A.6, 4.A.7, and 4.A.8, with the original results displayed in columns 1 and 3 of each table. The UTIP index provides shorter time coverage of little under 14 years and hence entails a larger dynamic panel bias and less reliable FE estimates. Despite the fact that this would bias coefficients upward, the point estimate for the aggregate (total) trade variable is smaller and less precisely estimated than with the newly constructed Theil index. Of the remaining variables, however, all of those estimated with a certain degree of precision are similar to the original results but slightly larger with the UTIP index. The lagged dependent variable shows a

Table 4.7: Robustness to the TFP index of technological change, table 4.2 results

	(1) Labor	(2) TFP	(3) Labor, constant sample	(4) TFP, constant sample	(5) Labor	(6) TFP	(7) Labor, constant sample	(8) TFP, constant sample
Tech(t-2)	0.244**	0.212**	0.449	0.861**	0.243**	0.212**	0.477	0.864**
	(0.0915)	(0.0974)	(0.470)	(0.353)	(0.0917)	(0.0982)	(0.476)	(0.357)
Totaltrade(t-1)	-0.000105	-0.000741	0.000959	0.000872				
	(0.000509)	(0.000947)	(0.00155)	(0.00154)				
Total_lt(t-1)					0.000747	-0.00310	-0.00302	-0.00273
					(0.00328)	(0.00361)	(0.00484)	(0.00453)
Total_mt(t-1)					-0.000814	-4.96e-05	0.00326	0.00316
					(0.00320)	(0.00393)	(0.00562)	(0.00579)
Total_ht(t-1)					-1.61e-05	-0.000604	0.000665	0.000508
					(0.000889)	(0.00119)	(0.00150)	(0.00143)
GDP	-0.0280	-0.00645	-0.210	-0.223	-0.0258	-0.0106	-0.217	-0.229
	(0.0865)	(0.121)	(0.216)	(0.213)	(0.0910)	(0.123)	(0.218)	(0.216)
Education	0.00790	-0.0164	-0.0504	-0.0353	0.00763	-0.0169	-0.0488	-0.0333
	(0.0347)	(0.0439)	(0.0508)	(0.0508)	(0.0353)	(0.0436)	(0.0555)	(0.0551)
ValAddAgri	-0.000368	-0.00312	0.000870	0.00100	-0.000188	-0.00420	-0.000800	-0.000394
	(0.00429)	(0.00547)	(0.00749)	(0.00722)	(0.00463)	(0.00631)	(0.00927)	(0.00891)
FDI(t-1)	0.00290	0.00647	0.00291	0.00357	0.00299	0.00670	0.00264	0.00317
	(0.00322)	(0.00629)	(0.00700)	(0.00711)	(0.00338)	(0.00691)	(0.00774)	(0.00776)
Observations	903	552	386	386	903	552	386	386
R^2	0.679	0.666	0.629	0.638	0.679	0.667	0.630	0.639
# of countries	58	37	33	33	58	37	33	33
Year FE	YES	YES	YES	YES	YES	YES	YES	YES

Notes. Robust standard errors in parentheses, ***p<0.01, **p<0.05, *p<0.1. The lagged dependent variable has been included in the estimation, but is omitted from the output.

slightly higher degree of inertia, which can be expected to still be downward biased. Notably, the technological change control variable is still strongly associated with the UTIP index, but is slightly smaller and less significant in some of the specifications. This fits well with the fact that the UTIP also uses other data sources to arrive at their index, which naturally cannot be expected to have any association with the technology measure that is intimately connected to the underlying data. The fact that there are no major changes in the results nevertheless lends support to the validity of the newly constructed Theil index.[18]

Finally, an extensive outlier analysis has been conducted, wherein single influential observations have been identified and deleted from the estimation sample. The overall results remain qualitatively and quantitatively unaffected.[19]

[18] The remaining, more disaggregated specifications have also been tested on the UTIP index. Results show no qualitative changes between the two measures apart from the already displayed loss in magnitude and significance when using the UTIP index (available upon request).

[19] Added-variable and partial-leverage plots, values of Cook's D, DFBETAs for the trade variables, and regression results with influential observations excluded are available upon request.

4.6 Conclusion

This paper has attempted to shed some light on the impact of trade on wage inequality in developing countries. It expands on the existing literature in four ways: First, by introducing a newly constructed measure of technological change into the empirical analysis, it addresses concerns of omitted variable bias. Second, it employs a comparative advantage-based country classification based on relative skill endowments, thereby incorporating previous findings in the literature which demonstrate the superiority of such a classification over the previously used income-based country categories. Third, it classifies trade flows according to their technology content, measured by the degree of human capital necessary to produce the goods. Lastly, a consistent version of the Theil index of inter-industry wage inequality is used which provides a longer and more consistent time coverage than existing measures.

Estimation results show that the coefficients on the trade variable are rather heterogeneous once relative endowments are taken into account and technology effects are separated from trade effects. Furthermore, their size increases substantially once this heterogeneity is accounted for, but standard errors remain too large to reach significance.

Introducing a new control variable of technological change, empirical findings demonstrate the need to control for this source of potential omitted variable bias, since in particular the export results change substantially with the inclusion of the variable. Some effects appear only when the variable is included, or disappear with its inclusion. In line with the previous findings in the literature on skill-biased technological change, the technological change variable itself is found to significantly and substantially increase wage inequality throughout all specifications. The fact that the medium technology export variables are the most sensitive to the inclusion of the technological change variable suggests that this is also where most of the technological progress seems to be taking place, in particular in the relatively more skill-endowed developing countries. This is in line with the South-South trade hypothesis, stating that this is where the medium-skill endowed country groups should have their comparative advantage.

Regarding technology transfer, the proposition made in the previous literature that trade to and from developed countries is disequalizing due to the introduction of skill-biased technological change can only partly be confirmed. No such effects are found for high-technology trade, neither through exports, nor through imports. In terms of medium technology, exports have positive coefficients and seem to be disequalizing in all but the countries in the lowest education quartile. This finding fits with the South-South trade story as well as the technology transfer hypotheses, in particular the absorptive capacity argument. It is difficult, however, to disentangle technology transfer from comparative-advantage, "trade"-based effects. As for the trade effects, results are generally in line with Heckscher-Ohlin theory for low-technology trade, where equalizing impacts are mostly found. Again, the disequalizing for medium-low technology trade is in line with the predictions of both the South-South trade and the technology transfer hypothesis and it is difficult to isolate these effects in the current set-up. More research is needed to investigate the exact magnitude of these effects vis--vis one another.

4.A Appendix

Table 4.A.1: Sample means of main variables

Country	no. of years	Theil index	Total imports (in bn. USD)	Total Exports (in bn. USD)	Education quartile (average)	Years of education[20]	Value added in agriculture	FDI (in bn. USD)	GDP (in mn. USD)
Argentina	37	0.053	10.86	10.08	3	8.2	6.35	5.82	299676
Bangladesh	29	0.04	1.87	0.89	1	2.8	29.39	0.09	124108
Bulgaria	16	0.068	8.15	0	3	9.6	10.8	3.54	70078
Bolivia	37	0.054	0.52	0.34	2.5	5.4	18.69	0.23	13899
Brazil	40	0.123	24.33	26.95	2	4.3	8.82	10.86	970292
Barbados	37	0.052	0.14	0.08	3	7.4	8.31	0.01	3362
Botswana	11	0.032	2.2	3.18	3	7.9	2.11	0.3	17000
Central African Rep.	28	0.051	0.04	0.03	1	0.9	39.49	0.01	1737
Chile	39	0.062	6.72	8.75	3	7.7	6.93	2.86	110627
China	20	0.091	242.3	320.19	2	6.6	14.78	59.1	5750470
Cte d'Ivoire	30	0.054	1.07	1.02	1	1.3	26.18	0.09	20154
Cameroon	31	0.097	0.58	0.35	1.3	2.9	27.58	0.09	19586
Congo	22	0.077	0.28	0.27	1.5	2.1	13.56	0.03	3303
Colombia	38	0.037	5.21	4.39	2	5.4	16	2.33	212890
Costa Rica	35	0.042	1.26	1.32	3	7.2	11.73	0.35	26220
Cyprus	40	0.026	1.74	0.53	3	8.4	6.61	0.56	10973
Dominican Rep.	29	0.072	0.34	0.21	2	3.8	19.15	0.05	21663
Ecuador	39	0.041	2.05	0.95	2.6	6.1	18.53	0.31	44533
Egypt	35	0.061	6.96	2.97	1.4	3.5	20.05	2.02	168531
Fiji	35	0.053	0.34	0.19	3	7.4	19.63	0.07	3147
Gabon	27	0.077	0.97	0.93	1.5	2.5	6.56	0.04	7098
Ghana	31	0.096	0.4	0.27	2	3.2	56.72	0.02	17474
Gambia	27	0.013	0.02	0.01	1	0.6	29.06	0	709

[20]Countries classified as "high education" are in the 4th education quartile at some point in time and used in the aggregation of the trade data in that year. They are: Albania, Argentina, Armenia, Australia, Austria, Belgium, Bulgaria, Belize, Barbados, Canada, Switzerland, Chile, Cuba, Cyprus, Czech Republic, Germany, Denmark, Spain, Estonia, Finland, Fiji, France, Great Britain, Greece, Hong Kong, Croatia, Hungary, Ireland, Iceland, Israel, Japan, Kazakhstan, Sri Lanka, Lithuania, Luxemburg, Latvia, Malta, The Netherlands, Norway, New Zealand, Poland, North Korea, Romania, Russia, Slovakia, Slovenia, Sweden, Tajikistan, Tonga, Trinidad and Tobago, Taiwan, and Ukraine.

Honduras	35	0.061	0.18	0.11	2	3.1	25.15	0.02	9474
Indonesia	40	0.082	41.28	58.98	1.7	5.4	13.99	6.04	812897
India	33	0.085	32.24	25.83	1	3.1	26.52	5.57	1616655
Iran	19	0.043	9.67	3.37	1.5	4.1	13.07	1.03	343572
Jordan	39	0.083	2.65	1.41	2.2	5.5	5.8	0.56	13979
Kenya	33	0.078	2.02	1.55	2	4.8	31.05	0.06	38707
Kyrgyzstan	15	0.286	0.47	0.38	3	9.1	36.53	0.07	10466
Lesotho	7	0.247	0.74	0.18	2	5.3	9.34	0.04	2348
Latvia	17	0.044	1.77	0.93	3	9.1	5.91	0.36	18994
Moldova	17	0.017	0.95	0.5	3	9.2	22.47	0.18	9179
Mexico	40	0.055	39.85	12.03	2.1	6.9	5.61	14.14	1032215
Malta	39	0.013	1.35	0.98	3	7.8	3.94	0.24	4788
Mongolia	12	0.068	0.47	0.3	3	8.1	26.26	0.12	6652
Mozambique	16	0.252	0.24	0.05	1	0.8	32.22	0.08	6228
Mauritius	32	0.054	1.56	0.79	2	6.1	9.59	0.09	10669
Malawi	37	0.095	0.18	0.09	1	2.2	40.58	0.01	5945
Malaysia	41	0.033	31.43	32.69	2.7	6.5	16.99	2.8	157322
Pakistan	26	0.075	6.67	5.86	1	3	25.52	0.69	261666
Panama	26	0.052	0.59	0.17	3	7.8	7.85	0.47	19867
Peru	36	0.276	6.23	6.98	3	7.7	8.2	3.11	142417
Philippines	38	0.055	11.92	10.11	3	6.9	21.53	0.81	187871
Poland	21	0.033	97.19	88.6	3	9.8	4.11	16.4	582449
Senegal	28	0.05	0.89	0.39	1.1	3.3	19.84	0.07	13001
Singapore	40	0.055	52.41	56.26	2.4	6.3	0.63	9.25	96904
El Salvador	35	0.06	1.18	0.82	2	5	13.95	0.3	5447
Syria	25	0.13	3.77	2.4	1.3	4.5	22.62	0.51	44857
Thailand	41	0.055	21.73	20.56	2	4.5	16.2	2.17	260430
Trinidad & Tobago	38	0.189	0.92	0.79	3	8.1	1.89	0.55	14581
Tunisia	39	0.166	4.07	2.69	1.4	3.4	15.7	0.47	41407
Turkey	40	0.057	25.16	16.42	2	4.3	20.08	2.67	496604
Tanzania	22	0.114	1.47	0.76	1	4.6	34.96	0.37	28737
Uganda	18	0.189	0.67	0.39	1	3.6	41.07	0.14	18444
Uruguay	35	0.055	1.8	1.62	3	7.6	10.17	0.42	29256
Venezuela	39	0.042	4.65	4.03	2.3	4.6	5.32	0.54	131013
Yemen	10	0.073	1.95	0.22	1	1.7	11.73	0.21	38228
South Africa	11	0.061	27.8	31.65	2	7.8	3.12	3.76	340185
Zambia	27	0.043	1.15	1.06	2	3	14.97	0.03	11764

Table 4.A.2: GMM estimates, table 4.4 (columns 2 and 4) results

	(1)	(2)	(3)	(4)
	Exports		Imports	
Dep. var.: ln(Theil)	No tech	Tech(-2)	No tech	Tech(-2)
ln(Theil)(t-1)	0.780***	0.811***	0.781***	0.838***
	(0.0365)	(0.236)	(0.0362)	(0.257)
Total_lt(t-1)	-0.00158	-0.00160	0.000712	0.000799
	(0.00147)	(0.00132)	(0.000498)	(0.000566)
Total_mt(t-1)	-0.00119	-0.000463	0.00206	0.00353
	(0.00352)	(0.00563)	(0.00864)	(0.00695)
Total_ht(t-1)	0.00337	0.00245	-0.00512	-0.00622
	(0.00448)	(0.00864)	(0.00583)	(0.00490)
Totalimp/exp(t-1)	-1.11e-05	0.000282	-0.000582	-0.000225
	(0.00147)	(0.00305)	(0.00199)	(0.00220)
GDP	-0.0254	-0.00756	-0.0247	-0.00280
	(0.0943)	(0.126)	(0.0903)	(0.118)
Education	0.0110	0.00646	0.00824	0.00643
	(0.0353)	(0.0346)	(0.0348)	(0.0317)
ValAddAgri	-0.00154	-0.00153	-0.00126	-0.00154
	(0.00474)	(0.00436)	(0.00463)	(0.00470)
FDI(t-1)	0.00365	0.00347	0.00435	0.00412
	(0.00329)	(0.00317)	(0.00347)	(0.00341)
Tech(t-2)	0.241**	0.246***	0.243**	0.242***
	(0.0916)	(0.0869)	(0.0914)	(0.0882)
Quartile2(LMEC)	0.0903	0.0851	0.0893	0.0747
	(0.0852)	(0.106)	(0.0844)	(0.103)
Quartile3(UMEC)	0.116	0.111	0.110	0.0929
	(0.0990)	(0.115)	(0.0989)	(0.114)
Observations	903	845	903	845
R^2	0.679		0.679	
# of countries	58	58	58	58
Year FE	YES	YES	YES	YES
# of instruments		52		51
Hansen Test		0.856		0.716
Sargan Test		0.753		0.745
AR(1)		0.0168		0.0196
AR(2)		0.163		0.164

Notes. Standard errors in parentheses as indicated in the top column; *** p<0.01, ** p<0.05, * p<0.1. Lags have been restricted to lengths 6-8 in column 2, and 6-7 in column 4. The depth of lag lengths has been guided by the misspecification tests, as well as achieving a realistic value for the lagged dependent variable, which should be between the OLS-estimate of 0.904 (0-.893) and the FE estimate of column 1 (3) for column 2 (4). Similar results emerge with varying lag lengths (results available upon request). LMEC = lower middle education country, UMEC = upper middle education country.

Table 4.A.3: GMM estimates, table 4.4 (columns 2, 4, and 6) results

	(1)	(2)	(3)	(4)	(5)	(6)
	Total trade		Exports		Imports	
Dep. var.: ln(Theil)	FE	GMM	FE	GMM	FE	GMM
ln(Theil)(t-1)	0.778***	0.803***	0.778***	0.844***	0.778***	0.819***
	(0.0364)	(0.263)	(0.0360)	(0.207)	(0.0365)	(0.272)
Total_ht(t-1)	0.0121	0.0102	-0.0146	-0.0112	-0.0383	-0.0305
	(0.0154)	(0.0273)	(0.0147)	(0.0151)	(0.0353)	(0.0571)
LMEC*ht(t-1)	-0.0140	-0.0119	-0.0132	-0.0182	0.0166	0.0131
	(0.0154)	(0.0295)	(0.0296)	(0.0352)	(0.0203)	(0.0284)
UMEC*ht(t-1)	-0.0120	-0.00991	0.00279	0.00955	-0.0165	-0.0127
	(0.0158)	(0.0297)	(0.0318)	(0.0411)	(0.0200)	(0.0300)
Total_mt(t-1)	-0.0196	-0.0186	0.0127	0.0187	-0.0166	-0.0128
	(0.0128)	(0.0224)	(0.0300)	(0.0382)	(0.0205)	(0.0300)
LMEC*mt(t-1)	0.0263*	0.0245	-0.0124	-0.00919	-0.0338	-0.0287
	(0.0137)	(0.0274)	(0.0178)	(0.0208)	(0.0252)	(0.0382)
UMEC*mt(t-1)	0.0170	0.0150	0.0401	0.00955	0.0299	0.0243
	(0.0130)	(0.0302)	(0.0259)	(0.0411)	(0.0269)	(0.0419)
Total_lt(t-1)	0.0178*	0.0177	0.0101	0.00435	0.0253	0.0187
	(0.00973)	(0.0128)	(0.0178)	(0.0272)	(0.0269)	(0.0482)
LMEC*lt(t-1)	-0.0229**	-0.0222	0.0181	0.0171	0.0359	0.0304
	(0.0109)	(0.0163)	(0.0125)	(0.0105)	(0.0304)	(0.0430)
UMEC*lt(t-1)	-0.0163	-0.0150	-0.0310**	-0.0283*	-0.0387	-0.0329
	(0.0111)	(0.0199)	(0.0153)	(0.0147)	(0.0328)	(0.0458)
Totalimp/exp(t-1)			-0.000270	-0.000460	0.000958	0.00106
			(0.00169)	(0.00147)	(0.000785)	(0.000984)
GDP	-0.0219	-0.00783	-0.0209	0.00973	-0.0313	-0.0139
	(0.0943)	(0.133)	(0.0967)	(0.102)	(0.0910)	(0.136)
Education	0.00764	0.00440	0.0114	0.00454	0.0101	0.00780
	(0.0364)	(0.0348)	(0.0361)	(0.0286)	(0.0373)	(0.0348)
ValAddAgri	-9.47e-06	7.53e-05	-0.000925	-0.00113	-0.000881	-0.00129
	(0.00474)	(0.00446)	(0.00498)	(0.00456)	(0.00462)	(0.00502)
FDI(t-1)	0.00322	0.00314	0.00354	0.00335	0.00430	0.00397
	(0.00363)	(0.00348)	(0.00360)	(0.00316)	(0.00382)	(0.00412)
Tech(t-2)	0.252***	0.256***	0.255***	0.261***	0.242**	0.243***
	(0.0937)	(0.0881)	(0.0927)	(0.0870)	(0.0936)	(0.0894)
Quartile2(LMEC)	0.129	0.124	0.0917	0.0745	0.122	0.106
	(0.0908)	(0.145)	(0.0936)	(0.102)	(0.0872)	(0.134)
Quartile3(UMEC)	0.141	0.132	0.0908	0.0728	0.156	0.136
	(0.113)	(0.170)	(0.113)	(0.117)	(0.108)	(0.166)
Observations	903	845	903	845	903	845
R^2	0.68		0.681		0.68	
# of countries	58	58	58	58	58	58
Year FE	YES	YES	YES	YES	YES	YES
Control variables	YES		YES		YES	
# of instruments		57		57		56
Hansen Test		0.852		0.635	*	
AR(1)		0.0267		0.00924		0.0253
AR(2)		0.171		0.152		0.169

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. Ht = high technology, mt = medium-low technology, lt = low technology. LMEC = lower middle education country, UMEC = upper middle education country. Lags have been restricted to lengths 6-8 in column 2, 7-8 in column 4, and 6 in column 6. The depth of lag lengths has been guided by the misspecification tests, as well as achieving a realistic value for the lagged dependent variable, which should be between the OLS- and the FE estimate. Other lag lengths yield similar results (available upon request).
*The Hansen test of overidentification is omitted for this equation since the instrument lag length is restricted to one lag, meaning that the model is exactly identified. No well-behaved model could be found with lag lengths deeper than 1. Results using alternative lag lengths and depths are available upon request.

Table 4.A.4: Robustness to the TFP index, table 4.3 (columns 2 and 4) results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Exports				Imports			
Dep. var.: ln(Theil)	Labor	TFP	Labor constant sample	TFP constant sample	Labor	TFP	Labor constant sample	TFP constant sample
Ht_exp(t-1)	-1.11e-05	0.00128	0.00556	0.00520	-0.000582	-0.00240	-0.00733	-0.00713
	(0.00147)	(0.00337)	(0.00447)	(0.00435)	(0.00199)	(0.00349)	(0.00600)	(0.00574)
Mt_exp(t-1)	0.00337	0.000352	0.00554	0.00464	-0.00512	0.00212	0.0120	0.0127
	(0.00448)	(0.00563)	(0.0107)	(0.0107)	(0.00583)	(0.00895)	(0.0114)	(0.0108)
Lt_exp(t-1)	-0.00119	0.00187	0.0130	0.0130	0.00206	-0.0126	-0.0280*	-0.0282**
	(0.00352)	(0.00808)	(0.0130)	(0.0124)	(0.00864)	(0.0103)	(0.0140)	(0.0129)
Totalimp/exp(t-1)	-0.00158	-0.00301	-0.00569	-0.00525	0.000712	0.000463	0.00527	0.00479
	(0.00147)	(0.00305)	(0.00549)	(0.00532)	(0.000498)	(0.00273)	(0.00479)	(0.00466)
GDP	-0.0254	-0.0126	-0.224	-0.234	-0.0247	-0.00867	-0.216	-0.226
	(0.0943)	(0.124)	(0.224)	(0.221)	(0.0903)	(0.125)	(0.224)	(0.221)
Education	0.0110	-0.0163	-0.0489	-0.0347	0.00824	-0.0168	-0.0426	-0.0267
	(0.0353)	(0.0433)	(0.0537)	(0.0540)	(0.0348)	(0.0427)	(0.0551)	(0.0545)
ValAddAgri	-0.00154	-0.00422	-0.00100	-0.000359	-0.00126	-0.00488	-0.00363	-0.00316
	(0.00474)	(0.00658)	(0.00973)	(0.00923)	(0.00463)	(0.00632)	(0.00906)	(0.00858)
FDI(t-1)	0.00365	0.00894	0.00848	0.00861	0.00435	0.00802	0.00487	0.00504
	(0.00329)	(0.00774)	(0.00855)	(0.00823)	(0.00347)	(0.00779)	(0.00858)	(0.00824)
Tech(t-2)	0.241**	0.210**	0.379	0.834**	0.243**	0.214**	0.529	0.874**
	(0.0916)	(0.0976)	(0.460)	(0.342)	(0.0914)	(0.0970)	(0.475)	(0.357)
Constant	-0.576	-0.721	1.134	1.166	-0.574	-0.748	1.065	1.102
	(0.864)	(1.111)	(2.032)	(1.987)	(0.827)	(1.121)	(2.013)	(1.966)
Observations	903	552	386	386	903	552	386	386
R^2	0.679	0.667	0.631	0.64	0.679	0.667	0.632	0.641
# of countries	58	37	33	33	58	37	33	33
Year FE	YES	YES	YES	YES	YES	YES	YES	YES

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. Ht = high technology, mt = medium-low technology, lt = low technology.

Table 4.A.5: Robustness to the TFP index, table 4.4 (columns 4 and 6) results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Exports				Imports			
Dep. var.: ln(Theil)	Labor	TFP	Labor constant sample	TFP constant sample	Labor	TFP	Labor constant sample	TFP constant sample
Tech(t-2)	0.255***	0.226**	0.559	0.874**	0.242**	0.214**	0.626	0.891**
	(0.0927)	(0.0965)	(0.488)	(0.357)	(0.0936)	(0.0988)	(0.454)	(0.362)
Total_ht(t-1)	-0.0132	0.0193	-0.0141	-0.0164	0.0166	0.0329	0.0216	0.0205
	(0.0296)	(0.0560)	(0.0655)	(0.0584)	(0.0203)	(0.0247)	(0.0239)	(0.0216)
LMEC*ht(t-1)	0.00279	-0.0419	-0.0116	-0.0105	-0.0165	-0.0360	-0.0337	-0.0334
	(0.0318)	(0.0609)	(0.0684)	(0.0604)	(0.0200)	(0.0283)	(0.0282)	(0.0261)
UMEC*ht(t-1)	0.0127	-0.0170	0.0295	0.0297	-0.0166	-0.0380	-0.0354	-0.0346
	(0.0300)	(0.0566)	(0.0689)	(0.0602)	(0.0205)	(0.0262)	(0.0268)	(0.0246)
Total_mt(t-1)	-0.0124	-0.0231	0.0403	0.0390	-0.0338	-0.0471	-0.0462	-0.0400
	(0.0178)	(0.0323)	(0.0434)	(0.0370)	(0.0252)	(0.0305)	(0.0304)	(0.0258)
LMEC*mt(t-1)	0.0401	0.0988*	0.123**	0.119**	0.0299	0.0538	0.0707*	0.0655*
	(0.0259)	(0.0565)	(0.0495)	(0.0458)	(0.0269)	(0.0338)	(0.0361)	(0.0339)
UMEC*mt(t-1)	0.0101	0.0209	-0.0461	-0.0420	0.0253	0.0558*	0.0774*	0.0736*
	(0.0178)	(0.0346)	(0.0496)	(0.0412)	(0.0269)	(0.0320)	(0.0412)	(0.0386)
Total_lt(t-1)	0.0181	0.0247	0.00804	0.00625	0.0359	0.0289	0.0296	0.0198
	(0.0125)	(0.0250)	(0.0203)	(0.0173)	(0.0304)	(0.0322)	(0.0333)	(0.0304)
LMEC*lt(t-1)	-0.0310**	-0.0473	-0.0362	-0.0318	-0.0387	-0.0482	-0.0739	-0.0603
	(0.0153)	(0.0300)	(0.0243)	(0.0208)	(0.0328)	(0.0364)	(0.0445)	(0.0414)
UMEC*lt(t-1)	-0.0146	-0.0241	0.0229	0.0228	-0.0383	-0.0648	-0.113*	-0.101*
	(0.0147)	(0.0245)	(0.0301)	(0.0272)	(0.0353)	(0.0398)	(0.0586)	(0.0534)
Totalimp/exp(t-1)	-0.000270	-0.00340	-0.0146*	-0.0137*	0.000958	0.00220	0.00875	0.00812
	(0.00169)	(0.00351)	(0.00831)	(0.00778)	(0.000785)	(0.00294)	(0.00629)	(0.00594)
GDP	-0.0209	0.0102	-0.157	-0.162	-0.0313	-0.0227	-0.204	-0.202
	(0.0967)	(0.127)	(0.225)	(0.223)	(0.0910)	(0.120)	(0.227)	(0.221)
Education	0.0114	-0.00929	-0.0397	-0.0268	0.0101	-0.0223	-0.0370	-0.0240
	(0.0361)	(0.0494)	(0.0571)	(0.0568)	(0.0373)	(0.0481)	(0.0608)	(0.0600)
ValAddAgri	-0.000925	-0.00383	-0.000461	6.14e-05	-0.000881	-0.00444	-0.00296	-0.00248
	(0.00498)	(0.00658)	(0.00963)	(0.00904)	(0.00462)	(0.00616)	(0.00898)	(0.00837)
FDI(t-1)	0.00354	0.00525	0.0112	0.0119	0.00430	0.00982	0.00882	0.00922
	(0.00360)	(0.00917)	(0.0119)	(0.0113)	(0.00382)	(0.00848)	(0.0100)	(0.00951)
Observations	903	552	386	386	903	552	386	386
R^2	0.681	0.67	0.641	0.65	0.68	0.67	0.639	0.648
# of countries	58	37	33	33	58	37	33	33
Year FE	YES	YES	YES	YES	YES	YES	YES	YES

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. The lagged dependent variable as well as the country group dummies have been included in the estimation, but omitted from the output to save space. Ht = high technology, mt = medium-low technology, lt = low technology. LMEC = lower middle education country, UMEC = upper middle education country.

Table 4.A.6: Robustness to the UTIP measure, table 4.2 (columns 2 and 4) results

	(1)	(2)	(3)	(4)
Dep. var.:	Theil	UTIP	Theil	UTIP
Lagged dep. var.	0.781***	0.787***	0.781***	0.784***
	(0.0365)	(0.0533)	(0.0368)	(0.0543)
Tech(t-2)	0.244**	0.199**	0.243**	0.199**
	(0.0915)	(0.0927)	(0.0917)	(0.0928)
Totaltrade(t-1)	-0.000105	-0.000742		
	(0.000509)	(0.000493)		
Total_lt(t-1)			0.000747	-0.00324
			(0.00328)	(0.00285)
Total_mt(t-1)			-0.000814	0.00269
			(0.00320)	(0.00335)
Total_ht(t-1)			-1.61e-05	-0.00154
			(0.000889)	(0.00106)
GDP	-0.0280	-0.0219	-0.0258	-0.0238
	(0.0865)	(0.0863)	(0.0910)	(0.0886)
Education	0.00790	-0.0227	0.00763	-0.0205
	(0.0347)	(0.0393)	(0.0353)	(0.0397)
FDI(t-1)	-0.000368	0.00115	-0.000188	0.00112
	(0.00429)	(0.00383)	(0.00463)	(0.00395)
ValAddAgri	0.00290	0.00559*	0.00299	0.00438
	(0.00322)	(0.00304)	(0.00338)	(0.00272)
Constant	-0.563	-0.571	-0.587	-0.568
	(0.780)	(0.817)	(0.837)	(0.841)
Observations	903	805	903	805
R^2	0.679	0.7	0.679	0.7
# of countries	58	58	58	58
Year FE	YES	YES	YES	YES

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Table 4.A.7: Robustness to the UTIP measure of inter-industry wage inequality, table 4.3 (columns 2 and 4) results

	(1)	(2)	(3)	(4)
Dep. var.:	Theil	UTIP	Theil	UTIP
Lagged dep. var	0.780***	0.784***	0.781***	0.783***
	(0.0365)	(0.0545)	(0.0362)	(0.0550)
Ht_exp(t-1)	-0.00119	-0.00255	0.00206	-0.00540
	(0.00352)	(0.00393)	(0.00864)	(0.00562)
Mt_exp(t-1)	0.00337	0.00350	-0.00512	0.00332
	(0.00448)	(0.00603)	(0.00583)	(0.00585)
Lt_exp(t-1)	-1.11e-05	-0.000972	-0.000582	-0.00235
	(0.00147)	(0.00199)	(0.00199)	(0.00209)
Totalimp(t-1)	-0.00158	-0.00162	0.000712	-0.000380
	(0.00147)	(0.00188)	(0.000498)	(0.000755)
Tech(t-2)	0.241**	0.198**	0.243**	0.199**
	(0.0916)	(0.0927)	(0.0914)	(0.0925)
GDP	-0.0254	-0.0240	-0.0247	-0.0173
	(0.0943)	(0.0915)	(0.0903)	(0.0885)
Education	0.0110	-0.0201	0.00824	-0.0223
	(0.0353)	(0.0403)	(0.0348)	(0.0394)
FDI(t-1)	0.00365	0.00514*	0.00435	0.00510*
	(0.00329)	(0.00282)	(0.00347)	(0.00286)
ValAddAgri	-0.00154	0.000493	-0.00126	0.00135
	(0.00474)	(0.00404)	(0.00463)	(0.00404)
Quartile2(LMEC)	0.0903	0.0710	0.0893	0.0711
	(0.0852)	(0.117)	(0.0844)	(0.118)
Quartile3(UMEC)	0.116	0.107	0.110	0.110
	(0.0990)	(0.131)	(0.0989)	(0.132)
Observations	903	805	903	805
R^2	0.679	0.7	0.679	0.7
# of countries	58	58	58	58
Year FE	YES	YES	YES	YES

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. Ht = high technology, mt = medium-low technology, lt = low technology. LMEC = lower middle education country, UMEC = upper middle education country.

Table 4.A.8: Robustness to the UTIP measure of inter-industry wage inequality, table 4.4 (columns 2 and 4) results

	(1)	(2)	(3)	(4)	(5)	(6)
	Total		Exports		Imports	
Dep. var.:	Theil	UTIP	Theil	UTIP	Theil	UTIP
Lagged dep. var.	0.778***	0.778***	0.778***	0.784***	0.778***	0.778***
	(0.0364)	(0.0562)	(0.0360)	(0.0555)	(0.0365)	(0.0564)
Total_ht(t-1)	0.0121	0.0200	-0.0132	0.0456	0.0166	0.0252
	(0.0154)	(0.0223)	(0.0296)	(0.0377)	(0.0203)	(0.0432)
LMEC*ht(t-1)	-0.0140	-0.0239	0.00279	-0.0537	-0.0165	-0.0288
	(0.0154)	(0.0224)	(0.0318)	(0.0393)	(0.0200)	(0.0430)
UMEC*ht(t-1)	-0.0120	-0.0223	0.0127	-0.0470	-0.0166	-0.0303
	(0.0158)	(0.0228)	(0.0300)	(0.0391)	(0.0205)	(0.0443)
Total_mt(t-1)	-0.0196	-0.0326	-0.0124	-0.0402*	-0.0338	-0.0474
	(0.0128)	(0.0197)	(0.0178)	(0.0224)	(0.0252)	(0.0414)
LMEC*mt(t-1)	0.0263*	0.0435**	0.0401	0.0626**	0.0299	0.0551
	(0.0137)	(0.0202)	(0.0259)	(0.0301)	(0.0269)	(0.0416)
UMEC*mt(t-1)	0.0170	0.0367*	0.0101	0.0417	0.0253	0.0590
	(0.0130)	(0.0208)	(0.0178)	(0.0257)	(0.0269)	(0.0441)
Total_lt(t-1)	0.0178*	0.0304**	0.0181	0.0299*	0.0359	0.0507
	(0.00973)	(0.0146)	(0.0125)	(0.0163)	(0.0304)	(0.0388)
LMEC*lt(t-1)	-0.0229**	-0.0404***	-0.0310**	-0.0421**	-0.0387	-0.0663
	(0.0109)	(0.0150)	(0.0153)	(0.0190)	(0.0328)	(0.0404)
UMEC*lt(t-1)	-0.0163	-0.0333**	-0.0146	-0.0306	-0.0383	-0.0583
	(0.0111)	(0.0149)	(0.0147)	(0.0186)	(0.0353)	(0.0396)
Totalimp/exp(t-1)			-0.000270	-0.00127	0.000958	-0.000162
			(0.00169)	(0.00255)	(0.000785)	(0.00126)
Tech(t-2)	0.252***	0.206**	0.255***	0.205**	0.242**	0.200**
	(0.0937)	(0.0971)	(0.0927)	(0.0964)	(0.0936)	(0.0965)
GDP	-0.0219	-0.0316	-0.0209	-0.0274	-0.0313	-0.0292
	(0.0943)	(0.0934)	(0.0967)	(0.0938)	(0.0910)	(0.0939)
Education	0.00764	-0.0191	0.0114	-0.0164	0.0101	-0.0215
	(0.0364)	(0.0422)	(0.0361)	(0.0410)	(0.0373)	(0.0428)
ValAddAgri	-9.47e-06	0.00122	-0.000925	0.000765	-0.000881	0.00150
	(0.00474)	(0.00417)	(0.00498)	(0.00425)	(0.00462)	(0.00424)
FDI(t-1)	0.00322	0.00398	0.00354	0.00522*	0.00430	0.00429
	(0.00363)	(0.00290)	(0.00360)	(0.00309)	(0.00382)	(0.00282)
Observations	903	805	903	805	903	805
R^2	0.68	0.703	0.681	0.702	0.68	0.702
# of countries	58	58	58	58	58	58
Year FE	YES	YES	YES	YES	YES	YES

Notes. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. Ht = high technology, mt = medium-low technology, lt = low technology. LMEC = lower middle education country, UMEC = upper middle education country.

4.B Appendix

Potential caveats of the sector-based approach for the measure of wage inequality

The following numerical example of three sectors, A, B, and C, demonstrates the potential caveats of the sector-based approach for the measure of wage inequality in the context of this paper. For reasons of simplicity, it is assumed that all sectors employ the same number of workers, which is stable over time. Furthermore, in the initial state before SBTC, skilled and unskilled workers earn the same wage, which is normalized to one and equal across sectors.

Table 4.B.1: Factor-biased SBTC, sector composition and average wage

		Sector A		Sector B			Sector C		
Wage growth of skilled workforce			20%		20%	40%		20%	80%
Composition of wages	Skilled	100	120	50	60	70	25	30	45
	Unskilled	100	100	150	150	150	175	175	175
Average wage		1	**1.1**	1	1.05	**1.1**	1	1.025	**1.1**

The first column in each sector then describes both the composition of the workforce and each group's total wage. SBTC raises the skill premium, leading to higher wages for the skilled. The second and third columns in each sector describe the resulting total wage for each skill group for different wage growth rates. With factor-biased SBTC only, the effect on the average wage depends on the composition of the workforce in each sector. The higher the share of skilled workers, the larger increase in the average wage. However, if factor-biased SBTC is asymmetrical (and thus also sector-biased), a larger increase in wages in one sector (e.g., 40 percent in sector B) can be partly or completely offset by the smaller share of skilled workers in that sector - which cannot be observed in the data at hand. One can see that in order to assess the overall effect of SBTC of wages, it is necessary to also take the distribution of wages within each sector into account. In the illustrated case, a between-sector measure would understate the effect of SBTC on the distribution of wages in the economy.

One might argue that the above reasoning also holds true for the opposite effect, namely trade-induced increase in the demand for unskilled labor. However, it is reasonable to assume that unskilled labor is more homogenous and exchangeable between sectors than skilled labor. Factor-biased SBTC favoring the unskilled therefore is therefore likely to affect unskilled wages rather symmetrically throughout the sectors of the economy. In sum, while there are a few caveats associated with employing a sector-based rather than a factor-based analysis, there is little reason to suspect that results will be distorted systematically. On the question of the importance of the within-group component of wage inequality, Conceição and Galbraith (2000: 71) argue that

"when the underlying data set is drawn from industrial classification schemes, the answer will generally be "not very important." Industrial classification schemes, after all, are designed to group together entities that are comprised of firms engaged in similar lines of work, and firms, like all bureaucracies, tend to maintain their internal relative pay structures comparatively stable from one period to the next."

When unskilled labor also (at least partly) profits from an increase in the wages of skilled labor within a sector, this mitigates the abovementioned problem of asymmetrical factor bias conflating the true extent of SBTC.

Bibliography

Acemoglu, D. (1998). Why do new technologies complement skills? Directed technical change and wage inequality. *Quarterly Journal of Economics*, pages 1055–1089.

Acemoglu, D. (2003). Patterns of skill premia. *Review of Economic Studies*, 70(2):199–230.

Alesina, A. and La Ferrara, E. (2005). Preferences for redistribution in the land of opportunities. *Journal of public Economics*, 89(5):897–931.

Alesina, A. and Rodrik, D. (1994). Distributive politics and economic growth. *The Quarterly Journal of Economics*, 109(2):465–490.

Anand, S. and Segal, P. (2008). What do we know about global income inequality? *Journal of Economic Literature*, 46(1):57–94.

Andridge, R. R. and Little, R. J. (2010). A review of hot deck imputation for survey non-response. *International Statistical Review*, 78(1):40–64.

Arellano, M. and Bond, S. (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.

Aristei, D. and Perugini, C. (2012). Inequality and reforms in transition countries. *Economic Systems*, 36(1):2–10.

Atkinson, A. B. and Brandolini, A. (2001). Promise and pitfalls in the use of "secondary" data-sets: Income inequality in OECD countries as a case study. *Journal of Economic Literature*, 39(3):771–799.

Atkinson, A. B. and Brandolini, A. (2009). On data: a case study of the evolution of income inequality across time and across countries. *Cambridge Journal of Economics*, 33(3):381–404.

Banerjee, A. V. and Duflo, E. (2003). Inequality and growth: What can the data say? *Journal of Economic Growth*, 8(3):267–299.

Barro, R. J. (2000). Inequality and Growth in a Panel of Countries. *Journal of Economic Growth*, 5(1):5–32.

Barro, R. J. and Lee, J. W. (2013a). A new data set of educational attainment in the world, 1950–2010. *Journal of Development Economics*, 104:184–198.

Barro, R. J. and Lee, J. W. (2013b). A new data set of educational attainment in the world, 1950–2010. *Journal of Development Economics*, 104:184–198.

Baumgarten, D. (2013). Exporters and the rise in wage inequality: Evidence from German linked employeremployee data. *Journal of International Economics*, 90(1):201–217.

Behrman, J. R., Birdsall, N., and Székely, M. (2007). Economic Policy Changes and Wage Differentials in Latin America. *Economic Development and Cultural Change*, 56(1).

Berman, E., Bound, J., and Machin, S. (1998). Implications of Skilled-Biased Technological Change. *Quarterly Journal of Economics*, 113(4):1245–1279.

Berman, E. and Machin, S. (2000). Skill-biased technology transfer around the world. *Oxford Review of Economic Policy*, 16(3):12–22.

Bernard, A. B. and Jensen, B. J. (1997). Exporters, skill upgrading, and the wage gap. *Journal of International Economics*, 42:3–31.

Bertola, G. (1993). Factor shares and savings in endogenous growth. *The American Economic Review*, 83(5):1184–1198.

Birdsall, N. (2008). Income Distribution: Effects on Growth and Development. In Dutt, A. K. and Ros, J., editors, *International Handbook of Development Economics, Volume 2*, chapter 48. Edward Elgar Publishing, Cheltenham, UK.

Blundell, R. and Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics*, 87(1):115–143.

Bourguignon, F. and Morrisson, C. (1990). Income distribution, development and foreign trade: A cross-sectional analysis. *European Economic Review*, 34(6):1113–1132.

Brosnan, S. F. and De Waal, F. B. (2003). Monkeys reject unequal pay. *Nature*, 425(6955):297–299.

Calderón, C. and Chong, A. (2001). External sector and income inequality in interdependent economies using a dynamic panel data approach. *Economics Letters*, 71(2):225–231.

Cappelen, A. W., Eichele, T., Hugdahl, K., Specht, K., Sørensen, E. Ø., and Tungodden, B. (2014). Equity theory and fair inequality: A neuroeconomic study. *Proceedings of the National Academy of Sciences*, 111(43):15368–15372.

Card, D. and DiNardo, J. E. (2002). Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles. *Journal of Labor Economics*, 20(4):733–783.

Chiu, W. H. and Madden, P. (1998). Burglary and income inequality. *Journal of Public Economics*, 69(1):123–141.

Clarke, G. R. G. (1995). More evidence on income distribution and growth. *Journal of Development Economics*, 47(2):403–427.

Coe, D. T., Helpman, E., and Hoffmaister, A. W. (1997). North-South R & D Spillovers. *The Economic Journal*, 107(440):134–149.

Cojocaru, A. (2011). Inequality and well-being in transition economies: A non-experimental test of inequality aversion. *ECINEQ Workin Paper No. 238.*

Conceição, P. and Ferreira, P. (2000). The young person's guide to the theil index: Suggesting intuitive interpretations and exploring analytical applications. *UTIP working paper No.14.*

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

Conte, A. and Vivarelli, M. (2007). Globalization and Employment: Imported Skill Biased Technological Change in Developing Countries. *IZA Discussion Paper No. 2797*, (2797):1–43.

Conte, A. and Vivarelli, M. (2011). Imported Skill-Biased Technological Change In Developing Countries. *The Developing Economies*, 49(1):36–65.

Cornia, G. and Kiiski, S. (2001). Trends in Income Distribution in the Post-World War II Period. *UNU-WIDER Discussion Paper No. 89.*

Cornia, G. A. (2014). *Falling inequality in Latin America: policy changes and lessons.* Oxford University Press, New York.

Davis, D. R. (1996). Trade Liberalization and Income Distribution. *NBER Working Paper No. 5693.*

De Santis, R. A. (2002). Wage inequality between and within groups: trade-induced or skill-bias technical change? Alternative age models for the UK. *Economic Modelling*, 19(5):725–746.

Deininger, K. and Olinto, P. (2000). Asset distribution, inequality, and growth. World Bank Policy Research Working Paper No. 2375.

Deininger, K. and Squire, L. (1996). A new data set measuring income inequality. *The World Bank Economic Review*, 10(3):565–591.

Deininger, K. and Squire, L. (1998). New ways of looking at old issues: inequality and growth. *Journal of development economics*, 57(2):259–287.

Dollar, D. and Kraay, A. (2002). Growth is good for the poor. *Journal of Economic Growth*, 7(3):195–225.

Dollar, D. and Kraay, A. (2004). Trade, growth, and poverty. *The Economic Journal*, 114(493):22–49.

Du Caju, P., Rycx, F., and Tojerow, I. (2012). Wage structure effects of international trade in a small open economy: The case of Belgium. *Review of World Economics*, 148(2):297–331.

Easterly, W. (2001). The middle class consensus and economic development. *Journal of Economic Growth*, 6(4):317–335.

Easterly, W. (2007). Inequality does cause underdevelopment: Insights from a new instrument. *Journal of Development Economics*, 84(2):755–776.

Edwards, S. (1997). Trade policy, growth, and income distribution. *The American Economic Review*, 87(2):205–210.

Egger, H., Egger, P., and Kreickemeier, U. (2013). Trade, wages, and profits. *European Economic Review*, 64:332–350.

Engelmann, D. and Strobel, M. (2004). Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *The American Economic Review*, 94(4):857–869.

Engerman, S. L. and Sokoloff, K. L. (1997). Factor endowments, institutions, and differential paths of growth among new world economies. In Haber, S., editor, *How Latin America Fell Behind*, pages 260–304. Stanford University Press, Stanford, CA.

Erhart, C. (2009). The effects of inequality on growth: a survey of the theoretical and empirical literature. *ECINEQ Working Paper No. 107.*

Essama-Nssah, B., Go, D. S., Kearney, M., Korman, V., Robinson, S., and Thierfelder, K. (2007). Economy-wide and distributional impacts of an oil price shock on the South African economy. *World Bank Policy Research Working Paper No. 4354*.

Ezcurra, R. (2007). Is income inequality harmful for regional growth? Evidence from the European Union. *Urban Studies*, 44(10):1953–1971.

Färe, R. and Primont, D. (2012). *Multi-output production and duality: theory and applications*. Kluwer Academic Publishers, Boston.

Feenstra, R. and Hanson, G. (2001). Global Production Sharing and Rising Inequality: A Survey of Trade and Wages. Technical Report 8372.

Feenstra, R. C. and Hanson, G. H. (1997). Foreign direct investment and relative wages: Evidence from Mexico's maquiladoras. *Journal of International Economics*, 42(3):371–393.

Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015). The next generation of the Penn World Table. *The American Economic Review*, 105(10):3150–3182.

Ferreira, F., Lakner, C., Lugo, M. A., and Özler, B. (2014). Inequality of opportunity and economic growth: A cross-country analysis.

Ferreira, F. H., Leite, P. G., and Wai-Poi, M. (2007). Trade liberalization, employment flows, and wage inequality in Brazil. Technical Report 4108.

Figini, P. and Gorg, H. (2011). Does Foreign Direct Investment Affect Wage Inequality? An Empirical Investigation. *The World Economy*, 34(9):1455–1475.

Foellmi, R. and Zweimüller, J. (2006). Income distribution and demand-induced innovations. *The Review of Economic Studies*, 73(4):941–960.

Forbes, K. J. (2000). A reassessment of the relationship between inequality and growth. *American Economic Review*, 90(4):869–887.

Forbes, K. J. (2001). Skill classification does matter: estimating the relationship between trade flows and wage inequality. *Journal of International Trade & Economic Development*, 10(2):175–209.

Galbraith, J. K., Choi, J., Halbach, B., Malinowska, A., and Zhang, W. (2015). A comparison of major world inequality data sets. *UTIP Working Paper No.69*.

Galbraith, J. K., Halbach, B., Malinowska, A., Shams, A., and Zhang, W. (2014). UTIP global inequality data sets 1963-2008: updates, revisions and quality checks. *UTIP Working Paper No.68*.

Galbraith, J. K., Jiaqing, L., and Darity Jr, W. A. (1999). Measuring the evolution of inequality in the global economy. *UTIP Working Paper No.7*.

Galbraith, J. K. and Kum, H. (2003). Inequality and Economic Growth: A Global View Based on Measures of Pay. *CESifo Economic Studies*, 49(4):527–556.

Galbraith, J. K. and Kum, H. (2005). Estimating the inequality of household incomes: a statistical approach to the creation of a dense and consistent global data set. *Review of Income and Wealth*, 51(1):115–143.

Galor, O. and Zeira, J. (1993). Income distribution and macroeconomics. *The review of economic studies*, 60(1):35–52.

Ghatak, M. and Jiang, N. N.-H. (2002). A simple model of inequality, occupational choice, and development. *Journal of Development Economics*, 69(1):205–226.

Glaeser, E., Scheinkman, J., and Shleifer, A. (2003). The injustice of inequality. *Journal of Monetary Economics*, 50(1):199–222.

Goldberg, P. K. and Pavcnik, N. (2007). Distributional Effects of Globalization in Developing Countries. *Journal of Economic Literature*, 45(1):39–82.

Goldin, C. and Katz, L. F. (1998). The origins of technology-skill complementarity. *The Quarterly Journal of Economics*, 113(3):693–732.

Gourdon, J. (2011). Wage inequality in developing countries: South-South trade matters. *International Review of Economics*, 58(4):359–383.

Gourdon, J., Maystre, N., and De Melo, J. (2008). Openness, inequality and poverty: Endowments matter. *Journal of International Trade and Economic Development*, 17(3):343–378.

Gruen, C. and Klasen, S. (2008). Growth, inequality, and welfare: comparisons across space and time. *Oxford Economic Papers*, 60(2):212–236.

Gruen, C. and Klasen, S. (2012). Has transition improved well-being ? *Economic Systems*, 36(1):11–30.

Harrison, A., McLaren, J., and McMillan, M. (2011). Recent Perspectives on Trade and Inequality. *Annual Review of Economics*, 3(1):261–289.

Haskel, J. E. and Slaughter, M. J. (2002). Does the sector bias of skill-biased technical change explain changing skill premia? *European Economic Review*, 46(10):1757–1783.

Heckscher, E. F. (1991). The effect of foreign trade on the distribution of income. In Heckscher, E. F. and Ohlin, B., editors, *Heckscher-Ohlin Trade Theory*, pages 39–69. MIT Press, London.

Ivaschenko, O. (2003). Growth and inequality: Evidence from transitional economies in the 1990s. In Eicher, T. and Turnovsky, S., editors, *Inequality and Growth: Theory and Policy Implications*, pages 155–198. MIT Press, Cambride, MA.

Jaumotte, F., Lall, S., and Papageorgiou, C. (2013). Rising Income Inequality: Technology, or Trade and Financial Globalization? *IMF Economic Review*, 61(2):271–309.

Jenkins, S. P. (2015). World income inequality databases: an assessment of WIID and SWIID. *The Journal of Economic Inequality*, 13(4):629–671.

Jones, R. W. (1971). A three-factor model in theory, trade and history. In Bhagwati, Jagdish N., K. C. P., editor, *Trade, Balance of Payments and Growth*, pages 3–21. North-Holland, Amsterdam.

Josten, S. D. (2003). Inequality, crime and economic growth. a classical argument for distributional equality. *International Tax and Public Finance*, 10(4):435–452.

Judson, R. A. and Owen, A. L. (1999). Estimating dynamic panel data models: a guide for macroeconomists. *Economics Letters*, 65(1):9–15.

Kaldor, N. (1955). Alternative theories of distribution. *The Review of Economic Studies*, 23(2):83–100.

Katz, L. F. and Autor, D. H. (1998). Changes in the wage structure and earnings inequality. In Ashenfelter, o. and Card, D. E., editors, *Handbook of Labour Economics*, pages 1463–1555. North-Holland, Amsterdam.

Keefer, P. and Knack, S. (2002). Polarization, politics and property rights: Links between inequality and growth. *Public Choice*, 111(1-2):127–154.

Klasen, S. and Gruen, C. (2001). Growth, income distribution and well-being in transition countries. *Economics of Transition*, 9(2).

Klein, M. W., Moser, C., and Urban, D. M. (2013). Exporting, skills and wage inequality. *Labour Economics*, 25:76–85.

Knowles, S. (2005). Inequality and economic growth: the empirical relationship reconsidered in the light of comparable data. *The Journal of Development Studies*, 41(1):135–159.

Köhler, M., Sperlich, S., and Vortmeyer, J. (2011). The Africa-Dummy in growth regressions. *Poverty, Equity and Growth Discussion Paper No. 94*.

Kremer, M. and Chen, D. (1999). Income-distribution dynamics with endogenous fertility. *The American Economic Review*, 89(2):155–160.

Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V., and Violante, G. L. (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica*, 68(5):1029–1053.

Kurokawa, Y. (2014). a Survey of Trade and Wage Inequality: Anomalies, Resolutions and New Trends. *Journal of Economic Surveys*, 28(1):169–193.

Lewis, W. (1954). Economic development with unlimited supplies of labour. *Manchester School of Economic and Social Studies*, 22(2):139–91.

Lipsey, R. and Carlaw, K. (2004). Total factor productivity and the measurement of technological change. *Canadian Journal of Economics*, 37(4):1118–1150.

López-Calva, L. F. and Lustig, N. (2010). Explaining the decline in inequality in latin america: Technological change, educational upgrading, and democracy. In López-Calva, L. F. and Lustig, N., editors, *Declining Inequality in Latin America: A Decade of Progress?*, pages 1–24. Brookings Institution Press and UNDP.

Loschky, A. (2010). Reviewing the nomenclature for high-technologythe sectoral approach. *Joint Research Centre, Scientific and Technical Report No. 57117, European Commission.*

Lundberg, M. and Squire, L. (2003). The simultaneous evolution of growth and inequality. *The Economic Journal*, 113(487):326–344.

Matsuyama, K. (2007). Beyond icebergs: towards a theory of biased globalization. *The Review of Economic Studies*, 74(1):237–253.

Meschi, E. and Vivarelli, M. (2009). Trade and Income Inequality in Developing Countries. *World Development*, 37(2):287–302.

Milanovic, B. and Squire, L. (2007). Does Tariff Liberalization Increase Wage Inequality? Some Empirical Evidence. In Harrison, A., editor, *Globalization and Poverty*, pages 143–181. University of Chicago Press, Chicago.

Molnar, T. S. (1972). *Utopia: The perennial heresy.* Tom Stacey Ltd.

Murphy, K. M., Shleifer, A., Vishny, R., et al. (1989). Income distribution, market size, and industrialization. *The Quarterly Journal of Economics*, 104(3):537–564.

Mussa, M. (1974). Tariffs and the distribution of income: The importance of factor specificity, substitutability, and intensity in the short and long run. *Journal of Political Economy*, 82(6):1191–1203.

Naranpanawa, A. and Bandara, J. S. (2012). Poverty and growth impacts of high oil prices: Evidence from Sri Lanka. *Energy Policy*, 45:102–111.

Neves, P. C. and Silva, S. M. T. (2014). Inequality and growth: Uncovering the main conclusions from the empirics. *Journal of Development Studies*, 50(1):1–21.

Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica*, 49(6):1417–1426.

Nunn, N. and Qian, N. (2014). Us food aid and civil conflict. *The American Economic Review*, 104(6):1630–1666.

O'Donnell, C. J. (2014). Econometric estimation of distance functions and associated measures of productivity and efficiency change. *Journal of Productivity Analysis*, 41(2):187–200.

Osberg, L., Smeeding, T. M., and Schwabish, J. (2004). Income distribution and public social expenditure: Theories, effects, and evidence. *Social Inequality*, pages 821–59.

Perotti, R. (1993). Political equilibrium, income distribution, and growth. *The Review of Economic Studies*, 60(4):755–776.

Perotti, R. (1996). Growth, income distribution, and democracy: what the data say. *Journal of Economic Growth*, 1(2):149–187.

Persson, T. and Tabellini, G. (1994). Is inequality harmful for growth? *The American Economic Review*, 84(3):600–621.

Piketty, T. (2014). *Capital in the Twenty-First Century*. Harvard University Press, Boston.

Pissarides, C. A. (1997). Learning by Trading and the Returns to Human Capital in Developing Countries. *World Bank Economic Review*, 11(1):17–32.

Robbins, D. J. (1996). Evidence on trade and wages in the developing world. *OECD Development Centre Working Paper No. 119*.

Roodman, D. (2009). A note on the theme of too many instruments. *Oxford Bulletin of Economics and Statistics*, 71(1):135–158.

Rubin, D. B. (2004). *Multiple imputation for nonresponse in surveys*. John Wiley & Sons, New York.

Savvides, A. (1998). Trade policy and income inequality: new evidence. *Economics Letters*, 61(3):365–372.

Schiff, M. and Wang, Y. (2004). On the quantity and quality of knowledge: The impact of openness and foreign research and development on north-north and north-south technology spillovers. *World Bank Policy Research Working Paper No. 3190.*

Solt, F. (2015). On the assessment and use of cross-national income inequality datasets. *The Journal of Economic Inequality*, (4):1569–1721.

Solt, F. (2016). The standardized world income inequality database. *Social Science Quarterly*, pages 1–15.

Spilimbergo, A., Londoño, J. L., and Székely, M. (1999). Income distribution, factor endowments, and trade openness. *Journal of Development Economics*, 59(1):77–101.

Stehrer, R. (2010). The effects of factor and sector biased technical change revisited. *Economic Change and Restructuring*, 43(1):65–94.

Stiglitz, J. (2014). Inequality is not inevitable. *New York Times.*

Stolper, W. F. and Samuelson, P. A. (1941). Protection and real wages. *The Review of Economic Studies*, 9(1):58–73.

Sukiassyan, G. (2007). Inequality and growth: What does the transition economy data say? *Journal of Comparative Economics*, 35(1):35–56.

Theil, H. (1967). *Economics and Information Theory.* North-Holland, Amsterdam.

Topel, R. H. (1997). Factor proportions and relative wages: the supply-side determinants of wage inequality. *The Journal of Economic Perspectives*, 11(2):55–74.

UNU-WIDER (2015). *World Income Inequality Database (WIID3.3).*

Viner, J. (1932). Cost curves and supply curves. *Zeitschrift für Nationalökonomie*, 3(1):23–46.

Vivarelli, M. (2014). Innovation, employment and skills in advanced and developing countries: A survey of economic literature. *Journal of Economic Issues*, 48(1):123–154.

Voitchovsky, S. (2009). Inequality and economic growth. In Nolan, B., Salverda, W., and Smeeding, T. M., editors, *The Oxford Handbook of Economic Inequality.* Oxford University Press, Oxford.

Windmeijer, F. (2005). A finite sample correction for the variance of linear efficient two-step GMM estimators. *Journal of Econometrics*, 126(1):25–51.

Wood, A. (1997). Openness and Wage Inequality in Developing Countries : The Latin American Challenge to East Asian Conventional Wisdom. *The World Bank Economic Review*, 11(1):33–57.

Wößmann, L. (2000). Specifying human capital: A review, some extensions, and development effects. *Kiel Institute for the World Economy (IfW) Working Paper No. 1007.*

Zhu, S. C. (2004). Trade, product cycles, and inequality within and between countries. *Canadian Journal of Economics*, 37(4):1042–1060.

Zhu, S. C. (2005). Can product cycles explain skill upgrading? *Journal of International Economics*, 66(1):131–155.

Zhu, S. C. and Trefler, D. (2005). Trade and inequality in developing countries: A general equilibrium analysis. *Journal of International Economics*, 65(1):21–48.

Eidesstattliche Erklärung

Versicherung gemäß §16 Prüfungs- und Studienordnung für den Promotionsstudiengang Wirtschaftswissenschaften

1. Die Gelegenheit zum vorliegenden Promotionsvorhaben ist mir nicht kommerziell vermittelt worden. Insbesondere habe ich keine Organisation eingeschaltet, die gegen Entgelt Betreuerinnen und Betreuer für die Anfertigung von Dissertationen sucht oder die mir obliegenden Pflichten hinsichtlich der Prüfungsleistungen für mich ganz oder teilweise erledigt.

2. Ich versichere, dass ich die eingereichte Dissertation "Economic Inequality: Causes, Consequences, and Measurement Issues. An Empirical Contribution." selbstständig und ohne unerlaubte Hilfsmittel verfasst habe; fremde Hilfe habe ich dazu weder unentgeltlich noch entgeltlich entgegengenommen und werde dies auch zukünftig so halten. Anderer als der von mir angegebenen Hilfsmittel und Schriften habe ich mich nicht bedient. Alle wörtlich oder sinngemäß den Schriften anderer Autoren entnommenen Stellen habe ich kenntlich gemacht.

3. Die Richtlinien zur Sicherung der guten wissenschaftlichen Praxis an der Universität Göttingen werden von mir beachtet.

4. Eine entsprechende Promotion wurde an keiner anderen Hochschule im In- oder Ausland beantragt; die eingereichte Dissertation oder Teile von ihr wurden nicht für ein anderes Promotionsvorhaben verwendet.

5. Des Weiteren ist mir bekannt, dass Unwahrhaftigkeiten hinsichtlich der vorstehenden Erklärung die Zulassung zur Promotion ausschließen bzw. später zum Verfahrensabbruch oder zur Rücknahme des erlangten Titels berechtigen.

Datum, Unterschrift