Essays in Poverty, Inequality and Political Economy

Dissertation

in order to acquire the doctoral degree
from the Faculty of Economic Sciences
at the Georg-August-Universität Göttingen

Submitted by

Rahul Lahoti

Born in Jalna, India

April 2017

Thesis Committee
First academic advisor   :   Prof. Stephan Klasen
Second academic advisor   :   Prof. Sanjay G. Reddy
Third academic advisor   :   Prof. Sebastian Vollmer

Acknowledgements

To Tulika and my parents

Contents

List of Tables

To focus on countries that might have a non-negligible effect on world economic dynamics or on the structure of the world-distribution, we isolate ones with a population of ten million or more. As a rough-and-ready criterion, we initially identify among these the top fifteen countries by per-capita income growth rate (see Table 3.7) according to surveys or national accounts. As can be seen, the countries that emerge from these two exercises are rather different. In particular, the survey based classification leads to some surprising results as it leads to the ......

List of Figures

1. Introduction

How to define development and measure it are strongly interlinked and have been discussed for several decades now. The conceptualization of development has evolved over time from a narrower understanding of the term and focusing on uni-dimensional measures to a broader understanding and using multi-dimensional measures. For a long time levels of Gross National Product (GNP) and growth in GNP were considered the best and sole measures of development. Sen (1988) argued that economic growth though an important contributor of development is insufficient as a measure and also that its distribution not just level needs to be considered before assessing its impact on welfare. Sen (1999) defined development as "removal of various unfreedoms that leave people with little choice and little opportunity of exercising their reasoned agency". This view of development as freedom was theorized by Sen with the Capability Approach. The Human Development Index developed jointly by Mahbub ul Haq and Amartya Sen (United Nations, 1990) was among the first broad-based measure to gain recognition that used the capability approach.

Individual welfare and its distribution across the society have garnered renewed interest recently in both academic and popular writings. This has been driven by moral concerns about increasing unequal distribution of resources and also by the need to understand the recent shifts across the world in opinion on global economic integration and its impact on inequality. Debates about the evolution of poverty, inequality and the extent to which the benefits of growth and development have been shared both within and across countries have all drawn attention. Increased availability of data and processing power have helped enhance our knowledge on level of inequities in the world. In the last decade, several researchers (Milanovic in particular) have used survey data from countries across the world to better understand global interpersonal inequality (see for e.g. Milanovic, 2012; Lakner & Milanovic 2015; Anand & Segal 2014; Edward & Sumner 2015). In the last few years the World Wealth and Income Database (WID) has amassed administrative data to discern distribution of income and wealth taking into account incomes of indiviudals at the top end of the distribution. The first two chapters of this thesis build upon this research to introduce the Global Consumption and Income Project (GCIP). GCIP is an effort to provide a deeper understanding of the evolution of material wellbeing both within and across countries, for regions and the world as a whole, extending from description to interpretation and explanation. GCIP provides a high standard of transparency and documentation, allowing for third-party replication, updating and the adoption of alternate assumptions for the selection and

treatment of data from the underlying universe, unlike current databases. GCIP also brings into focus several methodological issues with measuring distribution of income and consumption in a global context. The first chapter discusses the construction of the databases in the GCIP and the second chapter uses GCIP to study evolution of inequality and material well-being within and across 161 countries around the world from 1960 to 2015.

Multidimensional Poverty Index (MPI) developed by Oxford Poverty and Human Development Initiative (OPHI) and the United Nations Development Programme (UNDP) in 2010 (Alkire and Santos, 2010) is among the latest attempts to use capability approach to measure indiviudal welfare. MPI measures poverty at the household level by counting the number of deprivations that the household suffers. If the number of deprivations suffered by the household is above a cutoff (poverty line), then all individuals in the household are deemed poor. MPI measures deprivations in three dimensions, namely health, education and standard of living. The global MPI measures poverty in over 100 developing countries using household survey data. One of the drawbacks of the Global MPI is that it does not account for unequal distribution of resouces within the household as the measure is based on household-based analysis. The third chapter of this thesis uses data on individuals to define poverty measures that account for within household distribution of resources.

In addition to measuring development, it is critical to analyze how these outcomes can be improved. Political leaders play an important role in prioritizing sectors to focus government and policy attention on, framing policies and monitoring implementation of policies. Recent literature on India has shown that gender, religion, caste and other characteristics of the political leader play an important role in determining policy outcomes in the fields of education, health, economic growth and public infrastructure (Asher and Novosad, 2013; Bhalotra and Clots-Figueras, 2014; Bhalotra, Clots-Figueras, and Iyer, 2013). But the role of formal education of the politician in delivering development outcomes has not been investigated. In the final chapter of this thesis we study linkages between formal education of political leaders and development outcomes.

This thesis is an attempt to contribute to the literature on measuring welfare and how characteristics of political leaders impacts development outcomes. The main ideas, arguments and implications of each of the essays are condensed below.

Essay 1: The Global Consumption and Income Project (GCIP): An Overview

Growing concerns about unequal distribution of resources in the society and increased access to better quality data across countries have led to renewed interest in questions of absolute levels, gains and relative distribution of income and consumption across the world. Though there are several existing datasets that could be used to investigate these issues, most of them are severely limited. Almost all datasets combine the concepts of consumption/expenditure and income into one leading to incorrect conclusions on inequality comparisons across countries. Consumption and income not only have different levels for individuals but different distributions for the populations. These concepts are of independent interest as they represent concepts of advantage that are of evaluative concern for distinct reasons and because they provide different bases for empirical inference concerning material living standards. The level and distribution of the difference between the two (i.e. of savings or dissavings) may also be informative. Also several datasets that currently exist lack the transparency and flexibility to evaluate the impact of various assumptions that go into building any such databases.

In the first essay in this thesis (chapter 2), joint work with Sanjay Reddy and Arjun Jayadev, we introduce the Global Consumption and Income Project (GCIP), which has as its foundation the creation of two separate datasets (The Global Consumption Dataset (GCD) and The Global Income Dataset (GID)) containing a portrait of consumption and income of persons over time, within and across countries, around the world. The datasets present estimates of monthly real consumption and income of various quantiles of the population (a 'consumption/income profile') for the vast majority of countries in the world (161) for every year for more than half a century (1960-2015). The methodology of construction of the dataset allows for comparable data to be presented for an arbitrary number of quantiles (e.g. percentiles, ventiles, deciles, quintiles or other choices). The benchmark versions that are available for public use in beta versions report data in terms of mean levels of income and consumption by decile and in terms of 2005 and 2011 PPP dollars.

Construction of the dataset involves making various decisions and assumptions and can be broken down into four steps. In the first step, we collect data on relative distributions and mean levels for each country from various existing sources. In this step we are all inclusive, utilizing all sources for surveys like existing third-party databases, regional collection of surveys and

national statistical agencies. Where there is more than one survey for a country-year we select one based on a lexicographic ordering of various preferences that are clearly enunciated. In the second step, we standardize the distributions by converting all distributions that are not already in the required format (consumption or income distributions depending on the database) into estimated equivalents. The selected surveys for country-years consist of both consumption and income surveys. Where surveys of both kinds are available they differ, as the share of income tends to be higher for lower quantiles and the share of income lower for higher quantiles for income as compared to consumption distributions. Hence to make any meaningful comparison among distributions across and within countries and over time, we must transform the distributions. We derive a relationship between consumption and income distributions from country-years for which we have both kinds of surveys using a regression-based methodology and use this for the transformation. In the third step, where necessary we estimate a consumption mean for the GCD (Global Consumption Database) for survey-years where we have only an income mean and we estimate an income mean for the GID (Global Income Database) for survey-years where we have only a consumption mean so as to place the means too in more comparable units. We also attempt to detect means that are extreme outliers so as to enhance data reliability. Also we convert all means into common currency units in PPP terms so that levels are comparable over time and space. In the final step, using the mean and distributional data previously generated, we estimate a Lorenz curve for the survey years (using standard parametric methods). For the non-survey years we estimate the consumption/income profile by interpolation or extrapolation by using the appropriate per capita growth rate figures from the World Development Indicators (WDI)[1] to create a time-weighted average of the 'perspectives' on the estimation year that are associated with the nearest survey-years. This set of procedures gives rise to a complete time-space tableau covering the world between 1960 and 2015. We discuss details of each of these steps in the first essay.

The resulting nearly continuous portrait of the evolution of the world consumption and income pattern is unique in several ways. To avoid the potentially misleading effects of pooling consumption and income surveys together as exchangeable, we standardize the welfare concepts used. GCIP provides a considerably more comprehensive portrait by covering a larger group of countries (161), a longer time period (1960 to the present as opposed to the mid- or late 1980s to the present) and annual portraits of the world distribution, drawing on the nearest available surveys. In addition, in the GCIP, we employ purpose-built tools for aggregation that allow us to analyze evolution of material well-being for sets of individuals belonging to user-defined country

---

[1] World Development Indicators. Accessed Feb 1st, 2014. Retrieved from http://data.worldbank.org/data-catalog/world-development-indicators.

groups (e.g. regions of the world or countries sharing other specified traits) in a manner that is not possible with other existing databases. GCIP can be used to estimate mean and consumption and income profile for any given year and country or aggregate of countries. Synthetic populations[2] generated by GCIP can be used to estimate any poverty measure (headcount ratio, poverty gap ratio, FGT measure etc.), inequality measure (Gini coefficient, ratio of mean to median, Palma ratio, Theil index etc.) or measure of inclusiveness in growth and development (for example measures of how widely shared growth or pro-poor growth has been). We have also developed, and intend to provide publicly, in-built tools for filling in missing data, enhancing data reliability. GCIP also meets a high standard of transparency, which upon final release of the dataset would allow for third-party replication, modification and updating and the adoption of alternate assumptions for the selection and treatment of data from the underlying universe, unlike any of the current databases. All documentation and beta version of the data is available online at www.gcip.info.

Essay 2: Who Got What, When and How? A Fifty Year Overview from the Global Consumption and Income Project (GCIP)

In the second essay in this thesis (chapter 3), co-authored with Sanjay Reddy and Arjun Jayadev, we use the GCIP datasets, methods and the tools described in the first essay to undertake two separate but interlinked tasks. First, we provide a portrait of the global distribution of income and its constituent regional and national distributions and describe their evolution over time. Second, we describe the degree to which income or consumption levels of individuals have increased at various points in the distribution both within and across countries. We investigate the changes in distribution of income for the top 15 fastest growing countries in the world over the period of 1990-2013 both within country and also as a group using the aggregation module. This also helps us characterize their contribution to the overall change in the global distribution.

We have several important findings. First, global incomes have increased across the board over the last three decades. Median income doubled from 2005 PPP $73 to 2005 PPP $154 per month from 2000 to 2015 after much slower growth in the previous decade. Second, the global income

---

[2] For the GCIP we create synthetic populations that consists of 100 'persons', each representing a percentile in the distribution, but we can generate such a population of any size. Indeed, a separate concept that we employ is that of a 'model population' in which each representative individuals stands in for a certain number of persons (e.g. 10,000) from a certain country and segment of that country's income distribution.

distribution has become substantially more relatively equal in the last three decades. Most of these improvements came in the period after 2000 and can be attributed to inter-country rather than intra-country effects: the very rapid growth of China. China's population has experienced an enormous advance that has shaped the world distribution as a whole and others' relative positions. The rest of the world excluding China has witnessed increase in median incomes of only twenty six percent between 1980 and 2013 while China's median has increased by 1188 percent in the same period; and the income distribution for the rest of the world is still twin peaked. After 2005 developing countries other than China have seen higher growth rates and impacted the global distribution, but the sustainability of this growth is uncertain. Third, country experiences vary widely, and there is some evidence of 'inequality convergence' with previously more unequal countries becoming more equal over time and vice versa, although the majority of the world's population lived in countries with rising inequality. Also countries from Sub-Saharan Africa have the highest income inequality in the world upending the `stylized fact' that Latin American countries are among the most unequal in the world. Fourth, we find that despite significantly high growth in the fifteen fastest growing countries, the impact of Chinese growth dominates in changing the world income distribution and increase in inequality with fast growth is not the norm. Finally, we undertake a welfare analysis and find that while global 'welfare' estimated by any standard income-focused aggregative welfare function has increased substantially over the last fifty years most of the change has happened in the period since 2000. The rest of the world excluding China has seen no Pareto-improvements in the period 1970 to 2000.

Our results come with several important caveats. Global estimates are prone to several sources of uncertainty both intrinsic to the methods used and arising from the plausible use of different methods. Areas in which such uncertainties arise include survey design and sampling, PPP estimation, standardization methods used to compare income and consumption and others. Although they are known to exist, not all of the resulting probable errors can be quantified. Given the various assumptions that go into the estimates, we have done extensive sensitivity analysis to test robustness of our results. We provide estimates for various PPP base years, with and without standardization for consumption and income surveys, and by excluding certain countries from the estimates.

Essay 3: How Serious is the Neglect of Intra-household Measures to Multi-dimensional Poverty Measures?

6

Most gendered analysis of poverty compare either proportion of men and women who are deemed poor based on household based poverty measures or compare poverty rates by gender of the head of the household. Both of these approaches are incorrect and deeply misleading. Though the ultimate goal of poverty and inequality measures is to track welfare of individuals, most of these measures use the household as the unit of analysis. The poverty status of the household is determined based on a poverty line and uni or multi-dimensions indicators. All individuals in the household are assigned the same status as that of the household. This assumes equal or needs-based distribution within the household, which is inconsistent with the theoretical literature and empirical findings on intra-household bargaining.

A poor household might have individuals who are not deprived in any or most dimensions to be actually deemed as non-poor, and vice versa. The bias this generates in household-based multidimensional poverty assessments depends on how the thresholds for household poverty in a dimension are set, or how the individual-level data is used to create a household-level indicator. The deprivation thresholds can be defined in a *restrictive* way where the achievement of the worst-off member of the household has to be above the threshold for the household to be non-deprived. In these cases the deprivation rates among individuals are *estimated to be higher* by household measures as long as not all households are indeed equally deprived in that dimension. But deprivation thresholds could also be defined in an *expansive* way, where only the achievement of the best-off individual has to be above the threshold for the household to be non-deprived. In such cases, the deprivation rates among individuals are *estimated to be lower* by household measures if not all are as well off as the best-off. UNDP and OPHI's MPI use a mix of indicator threshold definitions – restrictive and expansive – so that the net bias of their neglect of intra-household inequality is not clear a priori.

In the third essay (chapter 4), coauthored with Stephan Klasen, we present a multi-dimensional poverty measure at the individual level that accounts for intra-household inequality across the entire population. Using data from India, we use this measure to estimate individual poverty and inequality as well as the size of the bias of household-based analyses.

Women and older individuals in India are far more deprived and poor than men and younger individuals. This simple fact is obscured and gender and generational differences are absent when measuring poverty and inequality using the standard household-based approach. In particular, the poverty rate of females is higher by 14 percentage points than men in our individual MPI measure but only 2 percentage points higher when using the household-based measure. The poverty rate among individuals aged fifty and over is higher by 46 percentage

points than among children aged between 7 and 18 years of age in the individual measure, compared to only 2 percentage points when using the household-based measure. Using a decomposable inequality measure, we find the contribution of intrahousehold inequality to the total inequality in the individual deprivation score inequality to be 30% and total inequality is also some 30% higher using the individual-based measure, while inequality among the poor is found to be 5% smaller using the individual measure. We also find that in over 60 percent of households the average deprivation level of women in the household is greater than the average deprivation level of men.

Our approach to individualize poverty measurement is only a first attempt in this direction and is hampered by insufficient data on individual well-being in standard household surveys; improved data would likely lead to even larger differentials in poverty by age and gender, at least in a country such as India.

Essay 4: Are Educated Leaders Better for Education? Evidence from India

Education of a leader has been used as a proxy for quality of the leader. The belief that formal education instills in political leaders the ability to be more competent and honest has motivated new laws in two states in India disqualifying anyone without minimum education requirements from contesting local elections. But there is almost no research investigating the link between formal education of the leader and competence in India. The few studies at the national level provide only inconclusive evidence.

In the final essay of this thesis (chapter 5), co-authored with Soham Sahoo, we test the hypothesis that educated politicians are more competent and deliver better development outcomes. In particular, we analyze if state level political leaders with a college degree are able to provide better education outcomes than leaders who have not completed college. Data on schooling outcomes (learning outcomes in reading, mathematics and English, school enrollment and attendance rates, school infrastructure and various grants for schools) from various data sources are combined with education information on state leaders who could influence education policy to analyze the impact of educated leaders. We create a dataset that combines data on learning outcomes for rural areas from the Annual Status of Education Report (ASER) from 2006 to 2014, enrollment outcome data from various National Sample Survey (NSS) rounds from 2007 to 2014, annual data on school infrastructure and funding from the census of schools

– DISE - from 2010 to 2014 and data on leader's education level from Association of Democratic Reform (ADR) from 2004 to 2014. This dataset allows us to analyze impact of political leaders education on education outcomes for 10 years, the entire span of time for which data on formal education of the leader is available. Our district-level analysis along with child level data enables us to hold a range of institutional and cultural factors constant, and yields statistical power to detect even small effects on education outcomes. To identify the causal impact of educated politicians, we use a fuzzy regression discontinuity design where the proportion of leaders who have completed college is instrumented with the proportion of college graduate leaders who won in close elections against a non-college graduate leader. This design is based on the assumption that winners in close elections are quasi-random. We extensively test the validity of this assumption.

In our analysis we find that educated politicians do not perform any better than less educated politicians. Education outcomes of districts which elect educated politicians is no better than in districts where less educated politicians are elected. The null result is precisely estimated and is consistent across OLS and 2SLS specifications. Our results are robust to a range of checks – different outcomes, specifications, different definitions of key parameters in the analysis. We also repeat our analysis for different groupings defined by district and individual characteristics – poverty, urbanization, education levels, and family characteristics – and in most cases we find no impact of political leaders on education outcomes.

There are several important caveats with our analysis; and it should be interpreted as a first step in understanding linkages between formal education of the leader and development outcomes in India. It might be the case that educated leaders focus on outcomes in other sectors; or national and local leaders whom we don't study have a positive impact; or influencing outcomes in education sector is difficult. In future research we plan to analyze relationship between educated leaders at various levels and other development outcomes.

Even with the various limitations, our findings suggest that educated leaders are not more competent than less educated leaders. This has implications on policy of disqualifying candidates based on their education and use of education as proxy for quality.

2. The Global Consumption and Income Project (GCIP): An Overview[3]

Abstract

We introduce two separate datasets (The Global Consumption Dataset (GCD) and The Global Income Dataset (GID)) making possible an unprecedented portrait of consumption and income of persons over time, within and across countries, around the world. The current benchmark version of the dataset presents estimates of monthly real consumption and income for every percentile of the population (a 'consumption/income profile') for more than 160 countries and more than half a century (1960-2015). We describe the construction of the datasets and demonstrate possible uses by presenting some sample results concerning the distribution of consumption, poverty and inequality in the world.

2.1. Introduction: Aims of the Project

Increases in mean per capita income are often used as an index of a society's economic development. However, it is a metric that is widely recognized to be quite insufficient. In recent years, public debate has been concerned with whether growth experiences are 'delivering' by enhancing well-being. Some recent work has focused on broadening the indicators which are used to assess social progress (see for example Stiglitz, Sen, & Fitoussi, 2010) while other work has been concerned with the highly unequal distribution of gains, whether accompanied by sizable improvements in the level of income and reductions in poverty (as in China) or by relative stagnation in the incomes of a considerable portion of the population (as in the United States). In the last two decades the increased availability of high-quality data has enabled researchers to provide an integrated portrait of inequalities within and between countries. Such

studies of inequality have, however, not generally been integrated with analyses of income growth.

We describe below an effort to create resources that can help address a range of questions, related to absolute levels, gains and relative distribution, by offering plausible estimates of the income and consumption enjoyed by different portions of the population within countries and in the world as a whole over a reasonably long time period. Specifically, we introduce the Global Consumption and Income Project (GCIP), which has as its foundation the creation of two separate datasets (The Global Consumption Dataset (GCD) and The Global Income Dataset (GID)) containing a portrait of consumption and income of persons over time, within and across countries, around the world. The project aims not only to construct but also to analyze these data in future work. The datasets present estimates of monthly real consumption and income of various quantiles of the population (a 'consumption/income profile') for the vast majority of countries in the world (more than 150) for every year for more than half a century (1960-2015). The methodology of construction of the dataset allows for comparable data to be presented for an arbitrary number of quantiles (e.g. percentiles, ventiles, deciles, quintiles or other choices). The benchmark versions that we intend to make initially available for public use will report data in terms of mean levels of income and consumption by decile and in terms of 2005 and 2011 PPP dollars.[4]

Using the GCIP one can estimate a Lorenz curve, mean and consumption and income profile for any given year and country or aggregate of countries. This enables us to create a synthetic population[5] from which any poverty measure (headcount ratio, poverty gap ratio, FGT measure etc.), inequality measure (Gini coefficient, ratio of mean to median, Palma ratio, Theil index etc.) or measure of inclusiveness in growth and development (for example measures of how widely shared growth or pro-poor growth has been) can be calculated.

The resulting nearly continuous portrait of the evolution of the world consumption and income pattern is unique. It goes beyond the Penn World Tables in presenting estimates of the distribution within countries and it goes beyond recent analyses of the world distribution both in greatly extending the period covered and in presenting estimates for every year as well as for both income and consumption. Whereas with rare exceptions (for example Lakner & Milanovic,

---

[4] The summary statistics and the methods for the databases (Version 0.1) that we report here reflect their versions as of March 14[th], 2015 and the secondary data for this version was downloaded on or before that date, The databases are, however, being continuously updated.
[5] For the GCIP we create synthetic populations that consists of 100 'persons', each representing a percentile in the distribution, but we can generate such a population of any size. Indeed, a separate concept that we employ is that of a 'model population' in which each representative individuals stands in for a certain number of persons (e.g. 10,000) from a certain country and segment of that country's income distribution.

2013) such databases and studies based upon them have focused on relative inequalities alone, we provide data on levels of consumption and income so as to enable assessment of level and distribution together, as is required for analyses of topics such as the inclusivity of growth and development. We have also developed, and intend to provide publicly, in-built tools for filling in missing data, enhancing data reliability, and creating portraits of aggregates of countries. Our intent is that the GCIP should meet a high standard of transparency, allowing for third-party replication, modification and updating and the adoption of alternate assumptions for the selection and treatment of data from the underlying universe, unlike any of the current databases. Among the benefits of such an approach is likely to be that the database can eventually be kept up-to-date through the involvement of multiple users, ensuring that it remains current. The fact that inferences often depend greatly on very detailed data choices makes such transparency indispensable[6].

Constructing the data set involves undertaking several decisions with regard to the selection of data as well as with regard to the manner in which estimates are generated for country-years in which no household survey was undertaken. Here we document the process of construction and specific choices concerning data in greater detail. Some of the other methods we have developed (e.g. for Lorenz curve estimation and aggregation) and software programs will be provided online at the project website (www.gcip.info). We briefly describe the methods we have employed in the construction of the benchmark version of the database and presents results for a few countries and aggregates. Extensions of the primary database (for instance involving quintiles or ventiles rather than deciles or different PPP concepts and base years) are created using analogous methods.

2.1.1. Comparison with Existing Databases

Estimates based on per-capita income of countries have been present since the 1950s and have been used to estimate global inequality (see for example Nurkse (1953) for an early estimate of the world income distribution on this basis, drawing on data collected by the League of Nations and the still nascent United Nations). Since the mid-1990s, when the Deininger and Squire dataset (Deininger & Squire, 1996) was released, economists have had data on the distribution of income across many countries, if often in summary form. This availability in turn has led to greater efforts to try and extend the data (for example, through the World Income Inequality

---

[6] For the example of the dependence of global poverty estimates on such choices, the implications of which are often obscured, see Reddy and Lahoti (2015)).

Database (WIID)[7] developed by WIDER, to 'harmonize' it by taking measures to ensure its greater comparability, as for example with the Standardized World Income Inequality Database (Solt 2009) and to extend the data backwards in time (see e.g. Pinkovskiy & Sala-i-Martin, 2009, which forms estimates for as early as 1970). The World Bank has been developing global poverty estimates on the basis of its own collection of data since the late 1970s, and the World Bank's Povcalnet database has been available to the general public since 2001 as a result of demands for greater data access and transparency. This institutional collection of data has also been the basis for the influential work of Milanovic (2002, 2005). [8]

Our work seeks to go beyond these earlier efforts in at least four ways. First, we construct estimates of both consumption and of income. It is well-known that consumption and income not only have different levels for individuals but different distributions for populations. They are moreover of independent interest, both because they represent concepts of advantage which are of evaluative concern for distinct reasons and because they provide different bases for empirical inference concerning material living standards. The level and distribution of the difference between the two (i.e. of savings or dissavings) may also be informative. We therefore create separate income and consumption estimates for each country-year observation and quantile in the database. Second, we aim to create a complete time-space tableau, interpolating where necessary in order to estimate mean level of income or consumption for every country and year as well as for distinct quantiles of the population. Third, we allow for the aggregation of estimates of the level and distribution of income for user-defined regions and groups of countries. This capability relies on our having previously created estimates that are aligned in time in a given year, through interpolation where necessary. This aspect of our effort therefore builds on the preceding one. We have developed our own software and methods to merge distributions for these user-defined aggregates, providing a flexible capability for researchers and policy analysts. Fourth, we aim to provide documentation of our methods and tools that is as complete as practicable so as to permit the adoption of alternate assumptions in order to construct other versions of the databases and to promote ongoing improvement of methods, tools and data through suitable engagement of specialists and the general public.

2.1.2. How does the GCIP compare to more recent efforts?

---

[7] World Income Inequality Database Version 3.3: https://www.wider.unu.edu/project/wiid-%E2%80%93-world-income-inequality-database
[8] Recently, the World Bank has made available a Global Consumption Database, which provides a detailed household-survey based picture of consumption patterns within countries, but this is available only for a very recent comparison year. (presently, 2010). See http://datatopics.worldbank.org/consumption/ .

Lakner & Milanovic (2013) build upon Milanovic (2005) and seek to describe the global income distribution between 1988 and 2008. They analyze the evolutions of levels of income as well as the distribution of income over time. They choose a few benchmark years and describe the change in the global distribution over the period using surveys based on observations at or near to those years. Whereas they pool income and consumption data without adjustment we employ a 'standardized' income concept (drawing on a broader universe of both consumption and income surveys and estimating income from consumption surveys or vice versa), and employ a much longer time series, in addition to the features of the project that allow for additional dimensions of flexibility, as mentioned above. Although we adopt this standardized approach because we believe it to enhance comparability, our data can also be used in 'pooled' fashion if desired, in keeping with their procedure and that adopted more recently by the World Bank.[9]

In another recent exercise Dykstra, Dykstra and Sandefur (2014) queried the Povcalnet database using automated methods to create a cumulative distribution of income or consumption (pooled together in that database) for a large number of survey-years (from each of 942 surveys spanning 127 countries over the period 1977 to 2012). The resulting database can (as with the GCIP) be used for diverse purposes, some of which would have been very difficult without downloading the data in this comprehensive way. The exercise highlights the difficulty in accessing even nominally public data for research and replication in view of the restrictive format in which it is often presented, the prevalence of poor documentation and the contrasting value of fully publicly accessible datasets. In creating an earlier version of the GCIP we undertook a very similar exercise. However, we abandoned that effort because (a) the computational effort for the exercise was very high and the cumulative distribution could simply be replicated for the entire distribution for as many points as desired, and more flexibly and transparently, by replicating the reported parametric regressions that underlay the data, (b) the Povcalnet database is largely confined to developing countries and to years from the early 1980s onwards and (c) there was no reason to privilege Povcalnet as a source of survey data even for developing countries, for which there are other sources of data too. The GCIP has been constructed to differ in key respects. The GCIP has wider area and time coverage (due to inclusion of surveys from other sources, largely secondary but sometimes primary), it incorporates a standardized welfare concept (consumption or income, with one estimated from the other where necessary) making within and cross-country comparisons more meaningful, it allows for the estimation of all measures for every year (not just the survey year or a reference year around which surveys are grouped), it provides tools for

---

[9] See Ferreira et al (2015) for details on the pooling method. There are questions however as to whether such pooling is sensible (Reddy and Lahoti (2015)) which is why we endeavor to separate consumption and income estimates.

15

creating user-defined composites of countries in any given year, it provides flexibility in choices as to how to construct and update the dataset, and in choosing specific estimation methods for the Lorenz curve (as opposed to accepting the version which happens to be chosen by Povcalnet,

Table 2.1: Comparison of Various Global Datasets

Database Features	Penn World Tables	WIID (Version 2.0C)	SWIID	Povcalnet	GrIP	GCIP
Coverage by type of country	Both developing and developed countries	Both developing and developed countries	Both developing and developed countries	Both developing and developed countries	Both developing and developed countries	Both developing and developed countries
Temporal coverage	1950-2011(not all countries)	1960-2008	1960-2005	1980-2012	1990-2012	1960-2012
Level of consumption/income	Both, based on national accounts	Only one of consumption or income, and not for all surveys.	Neither	Consumption or income only, based on surveys wherever possible	Consumption or income only, based on both NA and surveys	Both, based on surveys and national accounts.
Distribution by quantile	No	Yes	No	Yes, only survey years	Yes	Yes, all years
Adjustment of data on distributions to achieve greater comparability	N/A	No	Yes (through econometric estimation of Gini coefficients, adopting LIS as 'gold standard')	No	No	Yes (through econometrically estimated quintile-specific consumption-income ratio)
Interpolation for non-survey years	Yes	No	No	Not of reported quantiles or means, but implicit in reported poverty estimates	Yes- using distributions from surrounding region	Yes
Flexibility in modifying database according to alternate assumptions	No	No	In certain respects	In certain respects	In certain respects	Transparent about sources and methods so as to be flexible
Inequality measures	No	Gini only	Gini only	Selected, for survey years only	Potentially all	Potentially All

Aggregate over countries	Yes	No	No	Yes, but only for Poverty Measures	Yes for poverty	Yes, for poverty, inequality and the complete consumption or income profile (arbitrary number of quantile means)

which may reflect not only variable methods but sometimes generate invalid estimates of Lorenz curves). One of the key goals of GCIP is transparency, realized by providing documentation that is as complete as possible and access to all data and code to the extent feasible, in order to facilitate application of alternative assumptions in database creation or analysis. goals of GCIP is transparency, realized by providing documentation that is as complete as possible and access to all data and code to the extent feasible, in order to facilitate application of alternative assumptions in database creation or analysis.

Edward and Sumner (2013) have created a database closest in spirit and construction to ours. The Edward and Sumner GrIP ('Gr'owth, 'I'nequality and 'P'overty) model (version 1.0) takes distribution (quintile and decile) data and combines this with data on national population and on the mean consumption per capita in internationally comparable PPP $ to develop a database with similar aims to ours. However, the GCIP includes information before 1990, provides both consumption and income levels for each decile and allows for different PPP concepts as well as for market exchange rates. In this outline, we focus, however, on the present benchmark version which provides data in 2005 PPP dollars.

We do not attempt to discuss comprehensively the merits and demerits of previous efforts but instead seek to focus on the distinguishing features of the GCIP dataset. It is nevertheless useful to attempt to summarize the differences between our approach and existing efforts (see Table 2.1). We believe that the GCIP provides data for a wider set of countries, aggregates of countries, years and concepts, as well as tools for their analysis, than do other existing databases.

2.2. Construction of Global Consumption and Income Datasets

Constructing a consumption (or income) profile for a given country-year requires two distinct pieces of information: the relative distribution and the mean in that year. These two are sufficient

17

to create a unique profile of actual consumption (or income) levels of each decile in the country-year. We thus divide the process of creating the database into four distinct steps.

In the first step, we collect data on relative distributions and mean levels for each country from various existing sources. Where there is more than one survey for a country-year we select one, preferring consumption data sources for the consumption database and income data sources for the income database (Other choices are of course also possible, including to pool the income and consumption data without preferring one concept of advantage). Second, we standardize the distributions by converting all distributions that are not already in the required format (consumption or income distributions depending on the database) into estimated equivalents. The selected surveys for country-years consist of both consumption and income surveys. Where surveys of both kinds are available they differ, as the share of income tends to be higher for lower quantiles and the share of income lower for higher quantiles for income as compared to consumption distributions. Hence to make any meaningful comparison among distributions across and within countries and over time, we must transform the distributions. Although the conceptual case for doing so is strong this is rarely if ever done in international comparisons. In the third step, where necessary we estimate a consumption mean for the GCD (Global Consumption Database) for survey-years where we have only an income mean and we estimate an income mean for the GID (Global Income Database) for survey-years where we have only a consumption mean so as to place the means too in more comparable units. We also attempt to detect means that are extreme outliers so as to enhance data reliability. Fourthly, using the mean and distributional data previously generated, we estimate a Lorenz curve for the survey years (using standard parametric methods that have been found to perform acceptably in recovering underlying true distributions, although other methods are available in case these fail). Finally for non-survey years we estimate the consumption/income profile by interpolation or extrapolation by using the appropriate per capita growth rate figures from the World Development Indicators (WDI)[10] to create a time-weighted average of the 'perspectives' on the estimation year that are associated with the nearest survey-years. This set of procedures gives rise to a complete time-space tableau covering the world between 1960 and as near as we can come to the present. We describe each step in detail below.

2.2.1. Creating the Universe of Surveys

The GCIP draws data on relative distributions from diverse sources, such as the EU-SILC

---

[10] World Development Indicators. Accessed Feb 1st, 2014. Retrieved from http://data.worldbank.org/data-catalog/world-development-indicators.

database (for European countries), the LIS (previously the Luxembourg Income Study), the SEDLAC database (for Latin American countries), UNU-WIDER World Income Inequality Database (henceforth WIID), the World Bank's Povcalnet database, and Branko Milanovic's WYD database[11]. We are committed in principle to an ecumenical approach that integrates historical and contemporary data from *all* relevant sources, including country statistical offices, UN agencies, academic studies and private sector sources.[12]. Povcalnet is a collection of surveys starting from the early 1980s. Until recently, it covered only developing countries but now incorporates surveys for a number of developed countries, largely building on data from LIS. WIID is a collection of surveys from various secondary sources, covers both developed and developing countries and spans the period 1960-2012. Our third major source, the LIS, has harmonized data according to its chosen protocols from primary surveys for over 40 countries mostly from upper and middle-income countries. Although it provides data in household equivalence-scale adjusted form we extract the data we use from the underlying databases in per capita form.

Our first step is to generate a 'union' of all available distributional and level data for all the country-years of interest. The initial database thus constructed sometimes contains more than one observation for a country-year since multiple household surveys were undertaken in certain country-years and the same survey (in several instances with conflicting mean or distribution information) might be reported in multiple sources. The first task is therefore to refine the observations so as to arrive at one observation for each country and year. Surveys contained in GCIP may be reported as having a certain source, coverage of geographical area (national, or only urban areas), population and age, a certain assigned quality rating as stated in the underlying secondary source, concept of advantage (income vs. consumption, and specific income definition) and unit of analysis (household, individual, etc.). To choose one observation for country-years for which there are multiple we apply a lexicographic ordering to a set of selection criteria, which we discuss further below. The criteria and their sequence in the ordering are based on what we consider important considerations for common usage scenarios for the database. These can be altered if other usage scenarios are envisioned or indeed if users' judgments as to the relevance and importance of specific selection criteria differ from our own.

Before applying the various criteria, we restrict the universe of surveys to per capita surveys.

---

[11] www.lisdatacenter.org (accessed June 2015).
[12] GCIP also includes surveys for Cyprus, Hong Kong, Singapore and New Zealand from Branko Milanovic's World Income Database (WYD) as surveys for these countries were not available in the other secondary sources. We have also employed our own country research on specific individual cases to supplement our major sources, through correspondence with statistical agencies, identification of relevant historical documents etc. We list specific sources in our online appendix of country assumptions (see gcip.info).

This has the disadvantage of causing some loss of surveys and thus a reduction in the number of observations in our dataset, although much less than if we had chosen any other specific equivalence scale concept. For example, we are in the process of including the European Union Survey of Living conditions (EU-SILC) data. As this distributional data is reported an OECD-based equivalence scale we must recalculate the distributions in per-capita terms before including it. We prefer per-capita distributions for a number of reasons, and in keeping with the practice of other researchers (including the World Bank's Povcalnet, Milanovic and Lakner and others). Per capita surveys are simpler to analyze and to understand and correspond more directly to concepts in the national accounts. They are also the most common form of survey in the

Table 2.2: Welfare Concepts in GCIP Datasets

Global Consumption Database (GCD)

Welfare concept measured	Number of Surveys	Percentage
Consumption	830	42.7
Earnings, Gross	20	1
Earnings, Net	18	0.9
Factor Income	1	0.1
Income	285	14.6
Income, Disposable	588	30.2
Income, Gross	162	8.3
Monetary Income, Disposable	32	1.6
Monetary Income, Gross	8	0.4
Taxable Income, Gross	2	0.1
Total	1946	100

Global Income Database (GID)

Consumption	575	29.5
Earnings, Gross	20	1
Earnings, Net	18	0.9
Factor Income	1	0.1
Income	307	15.8
Income, Disposable	758	38.9
Income, Gross	217	11.1
Monetary Income	1	0.1
Monetary Income, Disposable	35	1.8
Monetary Income, Gross	12	0.6
Taxable Income, Gross	2	0.1
Total	1946	100

secondary data sources. The drawback of using only per-capita information is that differences in the real value of resources arising from variations in household size and composition are not taken registered. On the other hand, limiting our focus to per capita surveys greatly aids comparability. A variety of studies have shown that portraits of poverty, inequality, household

consumption behavior and other facts greatly depend on the equivalence scale chosen.[13] There is moreover reason to believe that even if the same equivalence-scales are being compared the extent and character of this dependence would vary greatly between country-years due to differences between country-years in the demographic composition of households belonging to different parts of the distribution. Whereas the exact nature of the dependence can be explored when the household level data is available, that is not possible when only summary results using a specific equivalence scale are reported, as is generally the case in the collections of data that we use. Rather than making our conclusions dependent in an unknown but very likely substantial way on the specific equivalence scales used we think it more sensible to use per-capita surveys. When it is reported that a survey uses an equivalence scale, typically insufficient detail is presented about the method that was used, making it difficult or impossible to compare distinct surveys meaningfully. As noted above, for LIS surveys, which report data using an equivalence scale, we obtain data in per capita terms using micro-data[14].

The lexicographic ordering of various criteria which we employ is as follows: whether a survey mean is reported, type of survey (consumption/income), the nature of the income/consumption definition, database source (e.g. EU-SILC, LIS, Povcalnet, SEDLAC, WIID, WYD, or primary source), area coverage, population coverage, quality as defined in the source database, source of the data as reported in the secondary database (e.g. source of a WIID observation) currency unit and survey series (as defined by statistical authority, e.g. German Socio-Economic Panel). As we are interested in both levels and distributions we prefer surveys with mean information over ones for which means are not reported. For the GCD, which focuses on consumption estimates, we prefer consumption surveys to income surveys (and vice-versa for the GID). Among income definition concepts we prefer concepts that are closer to arriving at total income net of taxes and transfers. The order of preference of income definition concepts appearing in the underlying databases (for which we draw upon the classification scheme and related definitions presented in the WIID) is as follows, from most preferred to least preferred: disposable income, disposable monetary income, gross income, gross monetary income, taxable disposable income, primary income, net earnings, gross earnings and finally a residual category for concepts that are not fully specified, i.e., we don't know if the reported data refers to net, gross or disposable income. Although it would be desirable in principle to make adjustments to the data based on relationships between the estimates corresponding to these distinct categories, in order to make them more comparable, we do not do so as we not have sufficient data corresponding to the

---

[13] See e.g. Buhmann et al (1988), Blaylock (1991), Coulter, Cowell and Jenkins (1992a and 1992b), Banks and Johnson (1994), Anand and Morduch (1996), Aaberg and Melby (1998), Cowell and Mercader-Prats (1999) or Sefil (2015) and the more recent literature cited therein.
[14] We use LIS's remote-execution system called LISSY to gain access to LIS's micro-datasets.

distinct concepts but the same countries or survey-years to establish these relationships.

Our order of preference of data by source employs the following ordering (earlier preferred to later): LIS, SEDLAC, EU-SILC, Povcalnet, WYD, WIID, primary source. This ordering reflects a number of judgments. Reported Povcalnet and LIS survey results are often compiled from primary data, while WIID is a collection of secondary data. We judge that Povcalnet and LIS may be more rigorously scrutinized and have a smaller probability of transcription or other errors as compared to WIID surveys and hence among global sources we prefer these two to the WIID. We view SEDLAC and the EU-SILC as being high quality sources of regional data (for Latin America and Europe respectively) and thus give high preference to them. Since LIS surveys have until recently included few if any developing countries and Povcalnet (as of recently) includes only select developed countries (corresponding to LIS countries) the overlap in terms of country-years covered by these is in proportional terms small. However, when there is an overlap we prefer LIS to Povcalnet for the reasons that LIS makes unit-level data available to us, and that LIS aims at achieving a higher degree of internal comparability among its surveys through specific effort at harmonization. The availability of unit-level data allows direct verification of the per-capita distributions calculated (which do in fact appear to coincide with the Povcalnet distributions for developed countries calculated from the same source). Due to this preference ordering, the external comparability of our estimates with Povcalnet based estimates for developing countries that derive to a larger extent from other sources, in particular, World Bank poverty estimates, is diminished (although we are to a large degree able to replicate these). The WYD overlaps heavily with Povcalnet but includes a few additional sources. To ensure greater comparability of GCIP with Povcalnet we place WYD after Povcalnet in our ordering.

We prefer surveys with broader area and population coverage and surveys deemed higher quality by the source database to others. WIID surveys report a quality rating but Povcalnet and LIS surveys do not report any quality rating. Given that Povcalnet and LIS are constructed using primary data and have stricter inclusion requirements we assign them the highest quality rating (but it must be remembered that this is only an ordinal characterization). Among sources in WIID (or in principle any other secondary source) survey data reported as originally from LIS or from the Deininger and Squire database are preferred over other sources.[15] We prefer surveys that report means in local currency units over those which are reported in other units because the method of conversion into international units by the source can often be non-transparent. We also prefer surveys in which the survey series is known over those for which it is missing. Even

---

[15] For earlier years in particular, WIID draws on other sources, such as Jain (1975).

after applying all of these criteria we find that some country-years have multiple surveys. At this stage we choose among these that survey which leads to the survey source being more compatible with the portrait presented by other years' observations for the same country (especially the nearest survey years for which data are available) or apply other criteria[16]. In certain instances, we exercise our judgment and drop certain surveys or prefer a survey to which the lexicographic ordering would not have led[17].

2.2.2. Standardizing the Distributions

Surveys vary widely by their focus (e.g. type of advantage, such as consumption or income), as well as details of their method (e.g. length of recall period, level of detail in surveys, whether unobserved costs and benefits are imputed (such as the value of rent for self-owned residences) and survey frame as well as timing) making comparability between countries difficult (For a discussion of this point see Smith, Dupriez and Troubat (2014)). Of particular interest to us is that the definition of income varies widely between surveys. Some report gross income, others after-tax income and others still wider or narrower categories, often with somewhat obscure definitions. Table 2.2 reports the various income/consumption concepts used in surveys included in the GCD and GID, along with their frequencies, adopting the classification used in the WIID.

As is well known, the distribution of consumption is expected to be less unequal than the distribution of income. Those concerned with estimating global inequality or poverty almost universally recognize this concern but do not generally correct for it[18] (Ferreira et al (2015); Lakner and Milanovic (2013)). Comparing measures of inequality or poverty across countries can therefore be highly misleading. Similarly, aggregating information for groups of countries to obtain, a measure of poverty or inequality, for say, Sub-Saharan Africa becomes difficult and results obtained from combining income and consumption based surveys may lead to misleading results.

---

[16] After applying the lexicographic ordering we observe multiple surveys for same country-year in four instances in the present version of the GCIP, which we resolve as follows. In the case of Barbados (1970) we use the survey that refers to the economically active population over one which covers only 'income recipients'. For China 1995, Brazil 1970 and Colombia 1964 we keep the survey that allows for a more consistent data series across the years for the country. This is an exercise of judgment and users might prefer the dropped surveys to be part of the database in which case they can make that choice.

[17] We modify our lexicographic ordering in the rare instances where there are known issues of comparability of the survey with other surveys for that country. One example of this is Indian consumption survey for 1999. The Indian consumption survey in 1999 used a shorter recall period of 7 days as opposed to usual practice of using a thirty-day recall period in other survey years for India making comparisons with other surveys difficult. Similarly, in the universe of surveys in the GCIP, Russia has consumption survey data reported by Povcalnet and income survey data from the LIS. The surveys from these sources for Russia exhibit vast differences in means information. Applying our lexicographic ordering we might have picked the LIS reported surveys in the instances when the both are present for the same survey-year. But in this case to maintain consistency of information across the time series and to keep as large a number of compatible observations as possible we choose to use Povcalnet surveys over those from LIS. We provide a list in our online appendix of country assumptions (see gcip.info) of all the cases in which we exercise our judgment over and above applying the rules described earlier.

[18] Deininger and Squire (1996), in the context of their dataset, suggest adding 6.6 Gini points to Gini coefficients based on consumption to obtain the corresponding income Gini coefficients based on average difference across their dataset.

Table 2.3: Summary Statistics for Surveys used in Standardization

	LIS	WIID	Total
# Surveys	87	117	204
# of Surveys by Regions			
Africa & Middle East	10	32	42
East Asia & Pacific	12	21	33
Europe	45	20	65
Latin America & Caribbean	18	19	37
Post Soviet	1	17	18
South Asia	1	8	9
# of Countries from each region			
Africa & Middle East	2	19	21
East Asia & Pacific	3	5	8
Europe	12	7	19
Latin America & Caribbean	4	6	10
Post Soviet	1	8	9
South Asia	1	3	4
# of Surveys before and after 1995			
1995 or Later	65	65	130
Before 1995	22	52	74
# Types of Income Surveys			
Income, Disposable	87	36	123
Income, Gross	0	81	81
# of Surveys by Income Grouping			
Low income	0	20	20
Lower middle income	4	44	48
Upper middle income	42	39	81
High income	41	14	55

One effort to address this issue is the work of Solt (2009) who makes the assumption (plausible at least for developed countries) that the LIS may be treated a 'gold standard' and then tries to adjust other surveys using a regression based method to estimate a 'standardized' summary measure of the distribution of income (the Gini coefficient that would be expected to result from counterfactual and missing LIS surveys) in other countries. His database is confined to measures of inequality. Niño-Zarazúa, Roope, and Tarp (2014) also estimate standardized consumption distributions, by adjusting the share of each consumption decile by the average difference between income and consumption decile shares for a set of country-years which had both type of surveys.

We take a different approach here. As it turns out there exist in the WIID and the LIS a total of

204 instances across 71 countries in which there is both consumption and an income survey reported by the same statistical agency for the same country-year. For most of these (more than ninety percent) information on consumption and income for the survey year was collected from the same survey. These survey countries are spread across all geographical regions of the world and across various country income groupings (Table 2.3).

We use this information to estimate the expected relationship between income and consumption. Our purpose was to identify a regression relationship between consumption and income for each quintile[19]. (We use quintiles rather than deciles in order to maximize the number of observations, as in earlier years often only quintile data is reported) Given that the errors across the five regressions might be correlated (and indeed, the Breusch-Pagan test suggested so), we employed a Seemingly Unrelated Regressions approach to estimate the relationship[20]. When we wish to estimate consumption shares from income shares the regression formula we use is:

$SC_{ij} = \alpha_i + \beta_i SI_{ij} + \gamma_i X + \varepsilon$ (1)

Where SC is the share of consumption of quintile i, SI is the share of income of quintile i, X refers to a set of controls for country income level, region, income concept used in the survey and time. Finally i and j are subscripts for country and quintile respectively.

When we wish to obtain the income share, we redo this exercise and reverse the regressor and regressand to obtain

$SI_{ij} = \alpha_i + \beta_i SC_{ij} + \gamma_i X + \varepsilon$ (2)

Table 2.4a and 2.4b provides the results of these regressions. In both sets of regressions the r-squared is moderately high, ranging from 0.47 to 0.76.

Table 2.5 provides an indication of the performance of this regression by reporting the results of an in-sample prediction analysis. The ability to predict consumption shares with a degree of reliability gives us confidence as to its general applicability.

---

[19] As noted earlier there are various income concepts collected by surveys and the choice to employ them might have affected estimates of mean levels and distributions. We do not standardize among these various income concepts in GCIP.

[20] We experimented with several different specifications and also used a (more theoretically appropriate) Dirichlet regression to estimate our equation .The Dirichlet distribution is the multivariate generalization of the beta distribution. Thanks to its properties, it is a convenient parameterization for compositional data. First, the dependent variables are restricted to the [0,1] interval. Second, it ensures that the shares sum up to unity. Hence, it is a valid distribution for estimating quantile shares, i.e. $SC_{ij} \sim D(\alpha_i)$. However, the results do not differ significantly between the two estimations. Moreover, the Dirichlet regression assumption that all shares are negatively correlated is violated in our data. We therefore use the more standard SUR approach. See Emener, 2015 for an analysis of the difference between the estimation techniques.

Table 2.4a: Seemingly Unrelated Regression (SUR) for converting Income distribution into equivalent consumption distribution for Global Consumption Database

	Quintile				
	1	2	3	4	5
Income Quintile Share	0.363***	0.334***	0.342***	0.317***	0.340***
	0.031	0.029	0.028	0.032	0.027
Income Grouping (Base: Low Income)					
Lower Middle Income	-0.661*	-0.481	-0.183	0.436	0.894
	0.353	0.403	0.402	0.38	1.381
Upper Middle Income	0.099	0.493	0.628	1.071**	-2.196
	0.396	0.453	0.452	0.426	1.553
High Income	0.144	0.513	0.723	1.180**	-2.817
	0.44	0.505	0.505	0.478	1.733
Region Grouping (Base: Africa & Middle East)					
East Asia & Pacific	2.276***	1.954***	1.371***	0.276	-5.929***
	0.316	0.363	0.363	0.348	1.25
Europe & Central Asia	1.219***	1.527***	1.422***	0.887**	-4.814***
	0.328	0.382	0.38	0.353	1.297
Latin America & Caribbean	-0.009	-0.13	-0.011	0.057	-0.015
	0.331	0.378	0.377	0.358	1.293
Post Soviet	0.710*	1.370***	1.496***	1.258***	-4.894***
	0.389	0.447	0.447	0.425	1.534
South Asia	2.511***	2.460***	1.981***	0.922*	-7.896***
	0.473	0.545	0.542	0.509	1.86
Welfare Concept (Base: Income Disposable)					
Income, Gross	0.227	0.406	0.533*	0.767***	-2.054**
	0.253	0.288	0.287	0.271	0.987
Time (Base: 1995 or Later)					
Pre 1995	-0.131	0.075	0.312	0.413**	-0.58
	0.191	0.218	0.217	0.205	0.747
Constant	3.773***	6.230***	8.739***	13.164***	34.561***
	0.408	0.492	0.539	0.706	2.277
R2	0.663	0.649	0.621	0.47	0.634
N	204				

We use these regression formulae to obtain a derived implied consumption distribution when one has only an income distribution available for a country and a derived implied income distribution when one only has information on the consumption distribution. We undertake this exercise for the whole dataset so that every country can be assigned an income and consumption distribution (at least one original and at most one derived) for every survey year.

However, prior to the final assignment we must make an adjustment for the adding-up constraint that the sum of percentage shares in the derived distribution must sum to one hundred. Typically,

one is left with income or consumption that is unaccounted for by the simple application of the regression coefficients, for the reason that the regressions were undertaken independently. The sum of shares might be above or below 100. We think it reasonable that the unaccounted for income may be added or subtracted (depending on the direction of the error in the total) proportionally equally across quintiles. This is admittedly only one possible choice: we could

Table 2.4b: Seemingly Unrelated Regression (SUR) for converting Consumption distribution into equivalent income distribution for Global Income Database

	1	2	3	4	5
Consumption Quintile Share	0.325***	0.318***	0.345***	0.293***	0.324***
	0.034	0.032	0.031	0.034	0.031
Income Grouping (Base: Low Income)					
Lower Middle Income	0.197	0.273	0.486	0.557	-1.491
	0.392	0.436	0.431	0.456	1.499
Upper Middle Income	1.430***	1.668***	1.684***	1.206**	-5.990***
	0.421	0.47	0.465	0.492	1.616
High Income	2.112***	2.836***	3.114***	2.921***	-10.873***
	0.455	0.508	0.503	0.533	1.75
Region Grouping (Base: Africa & Middle East)					
East Asia & Pacific	1.002***	1.821***	2.219***	2.810***	-7.846***
	0.361	0.396	0.389	0.405	1.359
Europe	2.286***	3.602***	3.556***	2.616***	-12.140***
	0.352	0.395	0.389	0.405	1.349
Latin America & Caribbean	-0.485	0.294	0.899**	1.539***	-2.197
	0.364	0.406	0.401	0.424	1.395
Post Soviet	1.830***	2.452***	2.733***	2.887***	-9.879***
	0.418	0.469	0.464	0.489	1.611
South Asia	1.965***	3.243***	3.308***	2.758***	-11.272***
	0.532	0.591	0.581	0.606	2.023
Time (Base: 1995 or Later)					
Pre 1995	0.32	-0.177	-0.484**	-0.583**	0.895
	0.209	0.233	0.231	0.244	0.801
Constant	0.475	2.553***	4.988***	11.068***	48.899***
	0.392	0.48	0.551	0.779	2.192
R-square	0.686	0.753	0.761	0.631	0.755
N	204				

apply another rule of apportionment. However, in the absence of compelling reasons to do otherwise, we think this a sound choice. Because of the independence of the quintile-specific regressions it is also possible that the derived implied consumption or income distribution might break the monotonicity restriction i.e. that the share assigned to a lower quantile might be greater than the share assigned to a higher quantile for the same country-year's estimated distribution. We check for this and in our preferred specification have not encountered any instances of this

issue[21]. However, if non-monotonicity were encountered there would be ways of addressing this.[22] An example of the application of this method is provided by Brazil in 1996. The GCIP has an income survey for Brazil for 1996, which we convert to an estimated "equivalent" consumption distribution based on our cross-country regression procedure. After application of the regression coefficients the sum of the shares of quintiles is 99.95. The deficit of 0.05 points is assigned proportionally to all the quintiles so that each quintile's share is increased by the same percentage. The shares at various stages of the process are shown in Table 2.6.

2.2.3. Standardizing the means

While there has been substantial interest among researchers in the variance between survey and national accounts means (see for example Deaton 2005), there has been little or no examination to the best of our knowledge of the variance *between* means from surveys carried out in the same year for a given country. Our initial examination suggests that the differences can be extremely wide. For example, Bolivia has two surveys in WIID for 1997 which report monetary income means that differ by 30 percentage points (414 vs. 538 Bolivianos per month). This in turn means that although our lexicographic ordering gives us a particular mean, a slightly different ordering might have led us to choose a dataset with a very different level of income or consumption. This problem will plague any attempt to choose surveys. The mean number of surveys per country-year is 2.95 and the country-years with more than one survey have on average 3.78 surveys, and only thirty percent of country-years have only one survey (although as we noted at the outset this can be due to the same survey being reported by multiple secondary sources).

Table 2.5: In-Sample Predictions of Income to Consumption Regressions

	Mean of Income Share	Mean of Consumption Share	Mean of Predicted Consumption Share	Difference between Actual Consumption and Predicted Consumption Shares	Difference between Consumption and Income Shares
1st Quintile	4.93	6.45	6.45	0.005	1.53
2nd Quintile	9.26	10.73	10.74	0.007	1.4
3rd Quintile	13.67	15.02	15.03	0.01	1.36
4th Quintile	20.49	21.43	21.45	0.014	0.95
5th Quintile	51.69	46.36	46.39	0.027	-5.32

---

[21] Users who would prefer not to do this standardization or replace our method with their own could, once the GCIP code is released, easily turn off or replace the standardization method.

[22] In particular, we would propose to apply a second regression to the independently estimated shares and assign estimates based on this regression to the quintiles, after adjusting to satisfy the adding-up constraint.

In future work, we hope to provide a more comprehensive examination of the issue of disparate survey means and their implication for such concerns as the global income or consumption distribution. For now, we simply note the problem and attempt to standardize the means for the surveys that our ordering leads us to. As noted before, the universe of surveys provides various definitions of income and consumption. Furthermore, these are often reported in non-comparable units (for example by providing the information in real or nominal terms, in local currency or international currency units, and for different time periods).

Table 2.6: Conversion of Income Survey Distribution for Brazil, 1996 to Equivalent Consumption Distribution

	Consumption Shares			
	Original Income Share	After Regression	After Adjustment for Adding-up Constraint	Difference in Equivalent Consumption and Income Shares
Q1	2.5	4.67	4.67	2.17
Q2	5.8	8.04	8.04	2.24
Q3	10.1	12.18	12.19	2.09
Q4	18.2	18.98	18.99	.80
Q5	63.4	56.07	56.10	-7.30
Total		99.95	100	
Gini	57.7		48.7	
Poverty Rate ($1.25 2005 PPP)	10.23		0.83	

Our next task is therefore to construct a consumption and income mean for every country-year in comparable units. In order to do this, we seek to generate an estimate of the consumption or income mean for each country-year for which we have an observation. Whenever an estimate of the mean was available from the survey with which we obtained the relative distribution, this was the preferred source of data[23]. This mean, usually expressed in local currency units (LCUs) of the survey year[24], was then converted to 2005 LCUs using local consumer price indices[25] wherever available (and in rare cases, where unavailable, the GDP deflator)[26].

---

[23] Lakner and Milanovic (2013) and World Bank's Povcalnet database also prefer survey means over national account means (see Anand and Segal (2008) for a discussion on the choice of means). Though we use survey means for our estimations, we will aim also to provide data on national account based means (GDP per capita and Household Final Consumption Expenditure (HFCE)) in the released version of GCIP. Mongolia is the only case for which we do not have any means from surveys and as a result we use means from national accounts as an alternative.

[24] We also attempt to adjust the means for any currency redenomination or change in currency that the country might have experienced. This is a non-trivial task as detailed historical knowledge of the country and its data sources is sometimes needed to do this.

[25] All our survey data is at national level and hence we use national CPI's unlike Povcalnet, which uses separate rural and urban survey components and inflation rates for India and China (Ferreira et al, 2015). It is not obvious whether using sector-specific data is superior because of the lack of uncontroversial inter-sectoral and sector-specific price indices, as discussed in Reddy and Lahoti (2015).

[26] Our source for inflation data is World Development Indicators (WDI). This contrasts with Povcalnet, which for some countries uses alternate CPI indices. For Taiwan, for which WDI does not maintain any data, we obtain data from Taiwan's National Statistical Office http://eng.stat.gov.tw/.

In order to make the estimates comparable across countries, we then converted them into common units by applying 2005 PPP exchange rates[27] and converting all data into monthly per capita units (for example if the survey estimate of consumption is for a weekly amount, we multiply it by 30/7). GCIP also includes country-specific conversion factors for other ICP PPP base years and other PPP concepts (e.g. PPPs for food) which could be used to obtain data in alternate PPP units and market exchange rates for all country-years. Note that in all cases we use the unitary country-wide PPP. This contrasts with, for example, the World Bank approach, which uses sectoral PPPs for urban and rural areas for India, Indonesia and China, based on back-of-the-envelope assumptions about likely inter-sectoral differences. A fuller discussion on the issues involved in deciding to use unitary or sectoral PPPs for these major countries and the impact of specific assumptions in the case of poverty estimates can be found in Reddy and Lahoti (2015)[28].

Outlier Detection

Despite our best attempts at selecting the data carefully, the survey mean data that we are left with contain outliers. These are means that are implausible *prima facie* given other existing data. In many cases we are unsure of the source of the discrepancies, especially in light of the fact that we draw extensively on secondary data. We identify outliers using two criteria described below. A survey mean that is identified as an outlier by *both* the criteria is marked as an outlier and adjusted according to a procedure we will describe.

To identify outliers, we first run a separate regression for each country to identify the time trend in survey means for that country. In this step, we regress the survey mean with respect to time (years elapsed since 1960). If the survey mean is above or below two studentized residuals from the regression line we mark it as a potential outlier. We find that about 8% of our observations are marked as potential outliers using this criterion. Applying this 'internal' criterion in isolation would mark cases in which a country's economy *actually* experienced sudden growth spurts or severe and sharp declines as outliers since a linear time trend may not be able to account for sudden transitions. To avoid this we impose a second 'external' condition, namely that the annualized survey mean growth rate is within certain bounds of the national accounts based

---

[27] We use 2005 EKS PPPs for `individual consumption expenditure household' concept obtained from the International Comparison Program (ICP) website. Even though we use 2005 PPP exchange rate for the benchmark version of the database, this is not because we necessarily prefer it to other exchange rates. The choice of exchange rate depends on the research question. We are aware that exchange rates used have a substantial impact on the levels and also on global and regional inequality and have presented some alternate estimates using 2011 PPPs and market exchange rates in Jayadev, Lahoti and Reddy (2015). In Reddy and Lahoti (2015) GCIP is used to calculate poverty estimates using 2011 PPPs and in Jayadev, Lahoti and Reddy (2015c) estimates of the evolution of the global middle class based on distinct concepts are explored using alternate exchange rates.

[28] Details of the PPP conversion factors and the ways in which they are implemented are available at www.gcip.info

growth rate in per capita gross domestic product. The acceptable band for the survey mean growth rate is defined by the growth rate of GDP per capita plus or minus twice the growth rate. (For instance, if the GDP per capita growth rate is 10% then the band is -10% to +30%). This criterion, while hardly restrictive, helps us to anchor the outlier detection criteria to a measure of validation external to the survey data, provided by the economy's growth rate. About sixty observations (5% of surveys with means data) are marked as outliers using both the criteria. Instead of completing discarding the outliers we view them as still providing relevant information and therefore instead adjust and retain them. The outlier means are adjusted (decreased or increased) up to the acceptable outer bounds of the time trend line. For example, outliers that are higher than the trend line are adjusted so that they have a value equal to the trend-line plus two studentized residuals. Our reasoning for doing so is that if we were to adjust the means to a higher level they would remain outliers according to our criteria, which would not serve the purpose of adjustment. At the same time, adjusting them to a level lower than the bounds would lead to treating outliers as requiring adjustment to a level lower than for reported survey means which are above the adjusted value of the mean but below the outlier detection bounds. It is hard to understand what would be the rationale for such a difference in treatment.

2.2.4. Generating a Lorenz Curve and Consumption/Income Profile

Having obtained or constructed means and distributional data for every survey year chosen, we estimate a Lorenz curve in parametric form using a widely employed regression framework (see Datt (1998); Miniou & Reddy (2009) for some discussion of the methods, also employed by Povcalnet). We prefer the generalized quadratic Lorenz curve estimation of Villasenor and Arnold (1989) for its theoretical properties but when the procedure fails to generate a valid estimated Lorenz curve we utilize the Beta Lorenz curve estimation due to Kakwani (1980) applied to quintiles[29]. When both of these methods fail to generate a valid Lorenz curve, which happens occasionally, we move to a third parametric approach due to Rasche et al (1980)[30]. If this were to fail (which it does not for *any* of the current distributions in the GCIP database) we would use a fourth parametric method due to Chotikapanich, D. (1993). Finally, in case all of these fail we create a piecewise linear consumption profile based upon 'connecting the dots' defined by the quantile means, following a method we have developed and tested (after which we can also calculate the associated Lorenz curve, which is strictly convex, as required for its

---

[29] In practice, when generating a valid Lorenz Curve, both procedures typically provide a reasonably good and similar fit to the data as captured by the sum of squared errors or other criteria. The Beta Lorenz curve fails the test of giving rise to a valid Lorenz Curve more often. The conditions of validity of these LC's are discussed in Datt (1998).

[30] To test the accuracy of Lorenz curves derived from the various parametric methods, we have compared income/consumption shares of various quintiles of the distribution obtained from parametric methods and that from unit-level data for a few LIS countries. Our initial findings indicate that all three methods perform very well in predicting the actual shares (within 1 percentage point in most cases). We hope to expand this analysis to all countries where we have access to unit-level data and report the results in a subsequent analysis.

validity). We chose to use the generalized quadratic Lorenz curve and the beta Lorenz curve in part because these are the parametric estimation methods used by the World Bank, and this would facilitate comparison of estimates, but one could equally use the Rasche method, which provides very similar results, based on our comparative examination of the methods for subsets of the data.

Once we arrive at an estimated Lorenz curve, we use it in combination with the estimate of the mean to generate a consumption profile consisting of an estimated mean income or consumption level for each decile of the country-year[31]. Specifically, the mean income of each decile is calculated by taking the share of total income accounted for by that decile, and multiplying it by the survey mean times the number of deciles (10). For example if the Lorenz ordinates for the first 2 deciles are 0.02 and 0.05 respectively and the mean income is $15, then the mean income of the first decile is $15*10*.02=$3, while the mean income of the second decile is $15*10*(.05-.02)=$4.5. We estimate decile means for the survey years in order to generate a lattice that can serve as a basis for interpolation and extrapolation of decile means to non-survey years. There are no deep-seated reasons for the use of decile means specifically for this purpose and we could have made a different choice.

Our goal is to estimate the consumption/income profile or set of quantile means for every country-year for the entire period covered by our database in order to obtain a complete 'consumption/income profile tableau'. In order to attempt to fill in the consumption/income profile tableau, we estimate the profile for intermediate years using growth rate figures from the World Development Indicators (WDI) or other sources where necessary, such as the Economist Intelligence Unit for most recent years, in order to interpolate or extrapolate consumption or income profiles for non-survey years. As noted below, the survey coverage is very limited before 1980. This is one reason why several researchers may have preferred to begin their empirical efforts after that date. Moreover, whether before or after that date they typically confine themselves to survey-year estimates, which may not be temporally aligned across countries, thus limiting the possibilities for comparison and aggregation across countries. We are contrastingly interested in trying to extend coverage as fully as possible in order to facilitate these tasks, not only by increasing the span of time covered but by filling in missing years. We fully recognize the concerns that such extension may raise, and accordingly try to do so according to carefully chosen assumptions. A larger amount of the data before 1980 is interpolated or extrapolated due to sparse survey information and thus has to be treated with

---

[31] In the case of the piecewise linear method for the estimation of the consumption profile, we would not need to generate a Lorenz curve at all.

greater caution. Users who would prefer to not use the interpolation/extrapolation techniques below we employ can always chose to restrict their analysis to only survey years[32], which are clearly marked in the GCIP.

There are two methods used to estimate a consumption/income profile for the non-survey year, viz.:

Extrapolation

If the non-survey year lies before or after the first/last survey year for which we have a consumption or income profile, then the consumption or income profile of that year is extrapolated (forward or backward) based on the survey year and the relevant per-capita growth rates. For example, if we want to estimate the consumption or income profile for a country and the last survey-year happens to be in a given prior year, then for the subsequent years, we extrapolate the consumption profile using the following formula iteratively:

$$M_t = M_{t-1} * (1 + g) \qquad (3)$$

where M is the estimated mean consumption/income of a decile, t is the year and g is the growth rate of mean consumption/income per capita between the two years.

Interpolation

If the non-survey year lies between two survey years for which we have a consumption or income profile, the consumption or income profile for this non-survey year is treated as a time-weighted average of the growth-adjusted consumption or income profiles (arrived at by extrapolating respectively backwards and forwards through applying the observed growth rates of mean per capita consumption or income) of the two survey years. This procedure is the same as described in Chen and Ravallion (2004) to impute means for non-survey years except that we extend the procedure to the overall distribution and estimate decile means in an analogous manner. Ferreira et al (2015) describe a procedure adopted by the World Bank more recently, in which the growth rates used in this process are adjusted by a ratio reflecting a presumed relationship between the growth rates from surveys and from national accounts in developing

---

[32] The estimates for survey years are not affected in any way by the interpolation/extrapolation. However, restricting the analysis to survey years would constrain cross-country comparison for a particular year as surveys are not generally lined up across countries. It would also limit the possibilities for aggregation if surveys for the countries being aggregated are missing for the year in question. Interpolation or extrapolation and aggregation are in this way intimately connected.

countries, and are thus lower than those in the national accounts. Specifically, they employ a multiplier of 0.87 to growth rates in all countries, except China and India where multiplier is 0.72 and 0.51 respectively, to make this adjustment[33]. This is unlikely to make very much difference for interpolation insofar as the adjustment applies to projections from both of the nearest survey years but it may make some difference to extrapolation. We do not employ such a further adjustment but it is would be a straightforward matter to construct a variant of GCIP that does so, either using the Bank's adjustment factors or another set of assumption (which for instance, distinguishes between the ratio applicable to countries of different kinds).

Since the consumption/income profiles for survey years are already expressed in comparable units ($2005 PPP in the benchmark version of the database) we therefore use the growth rates of real (inflation adjusted) per capita consumption to arrive at an estimated consumption or income profile for each non-survey year. Since growth rate information is available from different sources, we must establish a preference ordering for the growth rate data type and source, which is as follows, from most preferred to least: growth rate of household final per capita expenditure from the WDI, growth rate of 'per capita final consumption expenditure etc.' from the WDI, GDP per capita growth rate from the WDI, consumption per capita growth rate from the Penn World Tables, GDP per capita growth rate from the Penn World Tables, GDP per capita growth rate from the Total Economy Database (TED) (The Conference Board Total Economy Database 2010[34]) and finally GDP per capita growth rates from Angus Maddison 's historical statistics[35].

The earliest year to which we extrapolate our data backwards is 1960. This is because annual growth rates of mean consumption from national accounts for a wide variety of countries are available only starting then. There are some instances in which the growth rate data for the very earliest years is missing and we restrict the extrapolation to the first year when the data is available for these countries. The result in all of these cases is that there are gaps in the tableau. This not only affects the ability to define trends over the entire period but also to construct regional or global aggregates which are fully comparable over time. We hope to fill these gaps over time, in part by drawing on broad expert and public participation, or adopting other assumptions (such as extending trend growth rates backward or forward). In the meantime, one option is to discard from consideration those entities for which we do not have data over a sufficient period and another is to restrict the temporal scope of the analysis. For certain purposes, it may be tenable to compare alternatives which both do and do not contain certain

---

[33] Refer to footnote 48 in Ferreira et al. (2015) for more details on the adjustment factors used.
[34] Available at http://www.conference-board.org/data/economydatabase/
[35] The World Bank's Povcalnet also uses consumption per capita growth rates and GDP per capita growth rate data for interpolation/extrapolation. For more details refer to Footnote 48 in Ferreira et al (2015).

Table 2.7a: Summary Statistics for Surveys included in Global Consumption Database (GCD)

	1960-69	1970-79	1980-89	1990-99	2000-09	2010-13	Total
Number of Surveys	79	82	288	496	716	285	1946
% Consumption Surveys	14	12	24	48	54	41	43
% Surveys Covering Complete Population	56	68	78	98	100	100	93
% Surveys Covering all Areas in the Country	95	94	93	96	98	98	96
% Surveys with Mean Levels Information	41	49	65	90	96	100	86
Source of Surveys (%)							
EU-SILC	0	0	0	0	0	10	1
LIS	3	15	13	13	13	18	13
Povcalnet	0	0	15	35	61	59	42
Data from Statistical Office	0	0	0	0	0.4	0	0.1
SEDLAC	0	1	6	15	17	11	13
WYD	0	0	1	2	1	0	1
WIID	97	84	65	34	8	3	29
Regions (%)							
East Asia & Pacific	5	13	12	10	9	8	10
Europe & Central Asia	38	44	53	40	52	57	49
Latin America & Caribbean	23	20	19	28	22	20	23
Middle East & North Africa	5	6	3	4	3	2	3
North America	0	5	1	1	1	1	1
South Asia	19	7	4	3	3	2	4
Sub-Saharan Africa	10	5	8	14	11	9	11
Income Grouping (%)							
Low income	9	5	4	8	15	13	8
Lower middle income	31	16	23	23	13	7	22
Upper middle income	39	44	44	36	37	53	40
High income	21	34	28	33	35	27	30

countries, but one must be aware of the potential distortions that could arise as a result of specific countries dropping in or dropping out of the portraits of aggregates over time. The empirical examples of aggregates we provide in this paper do not include any adjustments for such non-comparability over time but that could be done in more careful subsequent work.

2.3.   Coverage of the Surveys

Tables 2.7a, 2.7b, 2.8a and 2.8b present summary statistics for the surveys in the GCD and GID. Tables 2.7a and 2.7b describe the number of surveys according to various criteria (e.g. source of

data, decade, region, income group), Tables 2.8a and 2.8b describe the number of countries in the databases according to these same criteria. Table 2.9 provides information on the density of the surveys36 by decades, region and

income group, i.e. the percentage of the total potential surveys (defined as the total number of country-years in the decade and region).

Table 2.7b: Summary Statistics for Surveys included in Global Income Database (GID)

	1960-69	1970-79	1980-89	1990-99	2000-09	2010-13	Total
Number of Surveys	79	82	288	496	716	285	1946
% Consumption Surveys	13	7	15	27	41	31	30
% Surveys Covering Complete Population	55	68	77	98	100	100	93
% Surveys Covering all Areas in the Country	95	94	93	95	97	98	96
% Surveys with Mean Levels Information	41	49	65	90	96	100	86
Source of Data (%)							
EU-SILC	0	0	0	0	0	10	2
LIS	2	15	13	14	14	18	14
Povcalnet	0	0	12	23	56	59	37
Data from Statistical Office	0	0	0	0	0.4	0	0.1
SEDLAC	0	1	6	17	18	10	13
WYD	0	0	1	2	1	0	1
WIID	97	84	67	44	12	4	34
Regions (%)							
East Asia & Pacific	5	13	12	10	9	8	10
Europe & Central Asia	38	44	53	40	52	57	49
Latin America & Caribbean	23	20	19	28	22	20	23
Middle East & North Africa	5	6	3	4	3	2	3
North America	0	5	1	1	1	1	1
South Asia	19	7	4	3	3	2	4
Sub-Saharan Africa	10	5	8	14	11	9	11
Income Grouping (%)							
Low income	9	5	4	8	15	13	8
Lower middle income	31	16	23	23	13	7	22
Upper middle income	39	44	44	36	37	53	40
High income	21	34	28	33	35	27	30

---

[36] The density of surveys is very similar in GCD and GID. This is because for a given country-year if there is only one survey available it will be used by both databases, and where there is more than one survey available a single survey will be chosen by each database (and hence the country-year coverage will be the same, even if the specific surveys selected are different). In Table 9 we therefore report the results for the GCD alone for simplicity.

There are a total of 1946 surveys in GCD over the fifty-five year period (1960-2015), from 161 countries. About forty three percent of the surveys are consumption surveys and more than ninety percent are nationally representative and cover the entire population. The coverage of surveys is sparse in the 1960s and 1970s with just over forty countries with surveys in each of these decades. The number of countries with at least one survey and the number of surveys with information on means both increase steadily in each decade, with rapid growth from the 1970s through the 1990s. Povcalnet is our biggest source of survey information, accounting for forty two percent of surveys in the GCD, followed by WIID (twenty nine percent), SEDLAC (thirteen percent) and LIS (thirteen percent). However, Povcalnet has almost no surveys in the first two decades, for which we instead rely heavily on WIID and to a lesser extent on LIS. As can be seen from Table 2.9, there are very few actual surveys available for earlier decades. For example, coverage is between zero and twenty five percent in the 1960s depending on the region. As a result, many estimates for years before 1980 in our databases are based on backwards extrapolation. When that is the case, users are cautioned to exercise judgment in their analyses.

Table 2.8a: Summary Statistics on Countries included in Global Consumption Database (GCD)

	1960-69	1970-79	1980-89	1990-99	2000-09	2010-13	Total
Number of Countries	44	44	98	133	151	111	161
Source of Data							
EU-SILC	0	0	0	0	0	9	29
LIS	2	6	18	22	17	17	45
Povcalnet	0	0	22	62	106	73	140
Primary Data	0	0	0	0	2	0	2
SEDLAC	0	0	5	12	14	7	23
WYD	0	0	3	4	2	0	5
WIID	42	38	50	33	10	5	116
Region							
East Asia & Pacific	2	6	11	17	20	11	21
Europe & Central Asia	16	14	44	48	49	44	50
Latin America & Caribbean	12	10	18	24	20	18	27
Middle East & North Africa	3	4	6	8	9	6	10
North America	0	2	2	2	2	2	2
South Asia	4	4	5	6	8	7	8
Sub-Saharan Africa	7	4	12	28	43	23	43

Additionally, one aspect of the paucity of coverage in earlier years is that there are several

countries for which the only distributional data we have are interpolated backward from later surveys. There are several examples of countries for which we do not have any or sufficiently reliable data for prior years. For a number of cases, e.g. Bahamas, Cuba, Germany[37], Israel, Kosovo, Puerto Rico, Somalia, former Soviet Republics, former Yugoslavia, Malta, Mongolia, Myanmar, West Bank and Gaza, etc. we have had to make specialized assumptions after undertaking research on the country's available data from historical records as well as economics analyses, to address issues related to splits and unifications, data gaps, or conflicting observations from distinct sources that are not otherwise resolved. The procedure we adopt to deal with cases of country splits and unifications is described in the appendix to this paper. The special assumptions made in the case of specific countries and years are identified in the online appendix of country assumptions (see gcip.info). We hope to try and acquire such information in future versions of the database, by eliciting the engagement of specialists and the general public. Although we have until now restricted ourselves to survey-based information, we intend in the future to explore the use of census data where feasible, although issues of comparability will have to be considered carefully.

Table 2.8b: Summary Statistics on Countries Included in Global Income Database (GID)

	1960-69	1970-79	1980-89	1990-99	2000-09	2010-13	Total
Number of Countries	44	44	98	133	151	111	161
Source of Data							
EU-SILC	0	0	0	0	0	1	29
LIS	1	5	17	21	26	34	45
Povcalnet	0	0	17	45	76	60	136
SEDLAC	0	0	3	11	16	9	22
WYD	0	0	2	4	3	0	5
WIID	42	39	59	51	27	6	119
Region							
East Asia & Pacific	2	6	11	17	19	11	20
Europe & Central Asia	16	14	44	48	48	44	50
Latin America & Caribbean	11	10	18	23	19	17	26
Middle East & North Africa	3	4	6	8	9	6	10
North America	0	2	2	2	2	2	2
South Asia		4	4	5	6	8	7
Sub-Saharan Africa	7	4	12	28	43	23	43

2.4.   Aggregation Module

---

[37] For example, for Germany prior to unification (1990), West Germany's distribution, mean and population information are used for Germany. We do not presently have any information on East Germany so prior to unification it is not included as part of Germany for these earlier years.

We have developed a module that can be used to obtain readily a consumption or income profile for an arbitrary grouping of countries. This can help to determine trends in poverty, inequality, growth in median consumption or income or other statistics of interest for any set of countries defined by region, income level, association membership or indeed any other criteria of interest[38]. These patterns can be juxtaposed with individual country experiences to understand how the set of countries is performing. We can perform various analytical exercises with data aggregated in this way such as decomposing contributions to levels (or changes in) inequality, poverty or other statistics into within-country and between-country components. The evolution

Table 2.9: Density of Surveys in Global Consumption Database (GCD) by Decade, Region and Income Group (% of all country-years that have survey information)

Region	1960-69	1970-79	1980-89	1990-99	2000-09	2010-13	1960-2013
East Asia & Pacific	2	6	19	25	29	23	17
Europe & Central Asia	7	9	18	36	55	29	29
Latin America & Caribbean	7	5	20	52	58	40	30
Middle East & North Africa	3	4	8	19	18	13	11
North America	0	20	15	35	30	38	21
South Asia	25	8	14	21	25	22	19
Sub-Saharan Africa	2	1	5	16	17	13	8
Income Group in 2010							
Low income	2	1	4	16	20	17	10
Lower middle income	6	3	11	29	38	23	18
Upper middle income	7	7	20	37	48	29	26
High income	5	8	18	33	39	26	22
Overall	5	5	14	30	37	24	19

of a group of countries can be surprising as it necessarily reflects the relative growth performance of different countries as well as their internal distributional dynamics. For instance, the evolution of inequality within a region (such as Latin America in recent years) may for this reason be different from what might be suggested by the evolution of inequality within individual countries. Several Latin American countries have experienced a dramatic decline in inequality in recent years: between 2000 and 2010 according to our estimates, the income Gini coefficient for Brazil has dropped 8 points (from 58 to 46). Chile's income Gini coefficient has dropped by 7 points (from 57 to 50). Despite the dramatic developments within a number of countries, our estimates indicate that in the same period the overall income Gini coefficient for Latin America and Caribbean has dropped too, but only by 4 points (from 55 to 51). This is because of the contribution of differential growth rates of different countries, which might not be

---

[38] Existing global datasets are not generally accompanied by an aggregation tool. In Povcalnet, users can obtain poverty headcounts for any arbitrary set of countries examined as a group but it does not provide a single aggregate distribution and hence cannot be used for inequality or inclusive growth analysis for the grouping of countries. Edward and Sumner (2013) present a notable exception.

apparent at first. This is an insight only made possible by looking at the composite of countries, as we are able to do. A few illustrative examples of applications of the aggregation module are discussed in Jayadev, Lahoti and Reddy (2015).

We confine ourselves here to briefly describing the method used to combine countries and to obtain a single consumption/income profile for the set of countries. We first obtain a consumption or income profile for all the individual countries within the grouping of countries and for a given year using the procedure described in previous sections. Next, employing a `poverty-line sweep' method, we obtain consumption levels for the 0.5 and 1.5 percentiles of the group. Specifically, we start at an arbitrary but low income/consumption level and calculate the percentage of population of each country that has income/consumption below this level. Then, using the population share of each country in the aggregate grouping we obtain the percentage of the group population at this level of the poverty line. We then iteratively adjust the level of the poverty line until we obtain the income/consumption level under which the desired percentage of the group population (i.e. 0.5 or 1.5) lies almost exactly, to a specified level of tolerance. Using these 0.5 and 1.5 percentile income/consumption levels for the group as starting points, we then raise the income/consumption level progressively in steps so as to obtain income/consumption levels at just over 100 points along the distribution, using continuous error corrections to adjust the size of the steps as we proceed so as to arrive at points within every or nearly every percentile interval. The resulting set of percentile points and the corresponding income/consumption levels are then connected linearly to obtain a consumption profile and to create a 'synthetic population', i.e. a representative population with the requisite profile, ordinarily consisting of one 'synthetic person' to represent the income or consumption level of each percentile. Using the synthetic population we can calculate any poverty or inequality measure, measure of inclusivity of growth etc. that we may wish to calculate for the group with complete flexibility[39].

2.5.   Comparability of Survey Data

Atkinson and Brandolini (2001, 2009) provide an account of problems of comparability of surveys across countries and discuss issues with existing databases. They emphasize the need for better documentation of the existing datasets, in order that there may be understanding of the methods used in constructing the datasets and greater clarity as to what kinds of comparisons are

---

[39] The computer code detailing the aggregation procedure will ultimately be provided on www. gcip.info.

meaningful and to what degree. However, they also place a considerable burden of responsibility on the researcher for understanding the data and for determining the extent of comparability. Some recent reviews of existing databases (e.g. Jenkins (2015); Smeeding and Latner (2015)) highlight the need for accessibility, replicability, transparency, quality of documentation and user-friendliness. One of the major motivations for developing GCIP was indeed to be able to provide transparent and replicable data. We plan to release the data, the underlying code and documentation of both methods used and of the various choices made (in particular through our websites www.gcip.info or www.globalconsumptionandincomeproject.org)

In addition, as already pointed out, GCIP is intended to provide a platform that is flexible. Researchers may choose to make different choices in data selection or processing and to generate their own version of the database (with suitable acknowledgment). Researchers could for instance easily change the exchange rate used for converting levels of income or consumption within countries to common currency units using the alternate conversion factors provided with the data or chose to use National Account means over survey means (both are also included). They could also easily modify the code we provide to turn off or even replace the procedure we employ to 'standardize' distributions or means, choosing to work only with income data, only with consumption data, or by pooling both. They may choose to change the order of priority we employ when selecting distributional or mean data for a given country-year or not to make such a selection at all and to work with the entire universe of data that we have collected. They may choose to work with survey year data alone, avoiding any interpolation or extrapolation to non-survey years.

In GCIP we at present select only per-capita data, standardize the welfare concept (by focusing on either consumption or income) and select surveys where multiple are available in order to improve quality and comparability according to our best judgment. We do not, and in several cases cannot, address other comparability issues due to a host of other reasons such as the use of diverse income concepts in different income surveys, varying survey designs across countries etc. However, in all instances we intend to provide the user with all the relevant details in accessible format to make the right decision for the purpose at hand.

2.6. Results[40]

---

[40] The results presented in this section are based on the datasets version dated Feb 1, 2016.

In this section, by way of conclusion we provide a few figures and tables that offer more specific indications of the kinds of analysis that are possible with the dataset. These are only illustrative examples. GCIP has been used to investigate poverty, inequality, inclusive growth, the profile of the middle class in Jayadev, Lahoti and Reddy (2015), Reddy and Lahoti (2015) and Jayadev, Lahoti and Reddy (2015b) and is being used to investigate these and other topics in a number of working papers under preparation. For the particular examples that follow we choose to focus on the global consumption distribution.

Figure 2.1: Superposition Graphs for World Consumption Distribution for 1980 and 2010
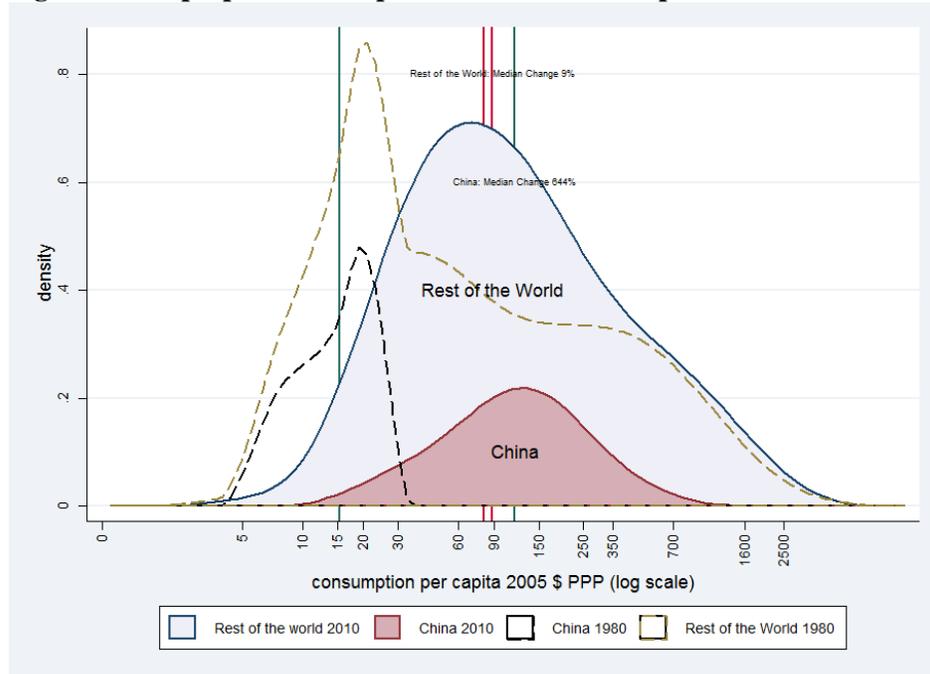
Source: Authors' calculations using GCD

2.6.1. Evolution of World Consumption Distribution

Figure 2.1 shows the evolution of world consumption distribution from 1980 to 2010 with the contribution of China and the rest of the world. It shows "twin peaks" in the 1980s (perhaps first identified by Quah (1996)). However, the period since then has seen the transformation of world consumption from a bimodal to a unimodal distribution and one in which the overall distribution has narrowed (also observed by Milanovic, 2005, Weisbrod et al, 2007 and noted in our own recent work, e.g. Jayadev, Lahoti and Reddy (2015a); see also Jayadev Lahoti, and Reddy (2015b) which use GCIP to describe the evolution of the global middle class for various understandings of the term.). This is primarily due to China's growth in the period and to some

extent the recent rapid growth in some other developing countries, especially after 2000. The Chinese consumption median increased by 644 percent between 1980-2010 while the rest of world experienced only a slight increase in the consumption median (nine percent). The corresponding change in means for the two groups is higher but a vast differential is here too apparent (885 percent vs. 26 percent). The world consumption distribution for 2010 might give a false impression of a symmetrical distribution of consumption across the world with little inequality but that is partly a result of using a log-scale for consumption expenditure. If we plot consumption instead of log-consumption the distribution seems highly skewed, as shown in Figure 2.2.

Figure 2.2: Kernel Density for World Consumption Distribution with a Non-log Scale
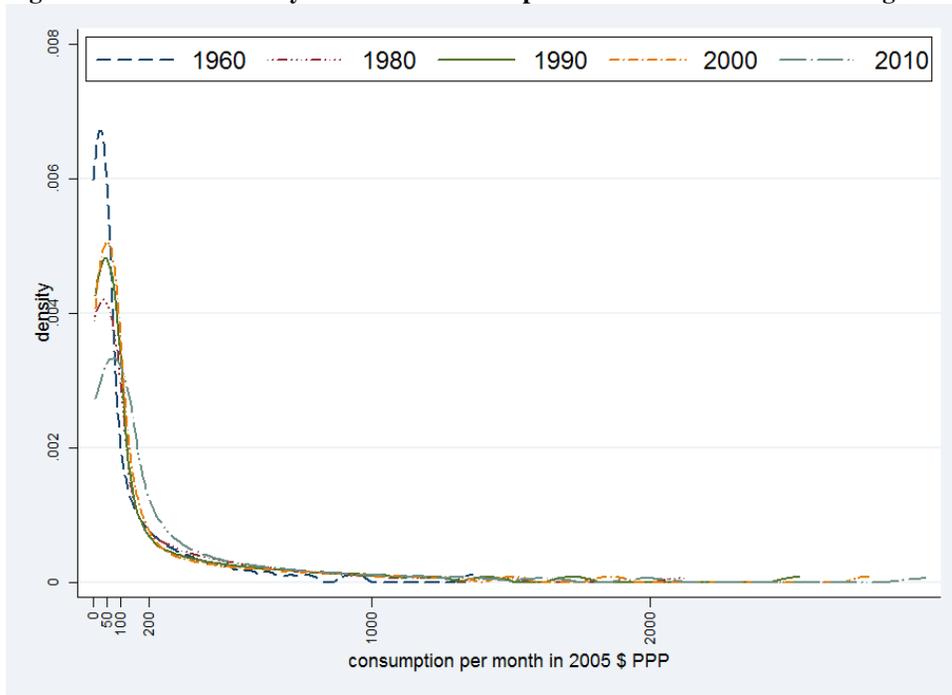
Source: Authors' calculations using GCD

The factors underlining the changing world distribution are underlined rather dramatically if one looks at the evolution of the global consumption Gini coefficient and the Theil index when including and excluding China, as in Figure 3. (This sort of analysis is straightforward when using the GCIP due its flexible aggregation feature). Rapid Chinese growth and its large population have meant that the global consumption Gini coefficient has fallen monotonically from its peak of .71 in the 1970s and 1980, to a low of 0.64 in 2013. However, excluding China from the picture suggests an altogether different picture. Consumption inequality for the world

without China was increasing from 1970 to 2000 and has since declined slightly. Over the period of 1970 to 2013 it has increased slightly from .64 in 1970 to .66 in 2013. Given the multitude of sources of uncertainty affecting these aggregate estimates – sampling, measurement and coverage errors in surveys, widely varying and inconsistent survey concepts and methods, difficulties with estimating and interpreting PPPs, poorly measured growth rates and other sources – it's not clear if these modest changes should be viewed as statistically or economically significant (see more on the sources of errors in global inequality estimates in Anand and Segal (2008)). It is interesting to note that the world excluding China has become more unequal (after 2000-2005) than the world including China, whereas the opposite was true earlier. China's inclusion has clearly also had a significant moderating influence on world inequality, in particular in the period 1980-2000. This was presumably because rapid Chinese growth served to reduce average gaps between the Chinese people and those richer than them. This effect overwhelmed the increase in inequality within the country and indeed increases in inequality elsewhere.

Figure 2.3: Inequality for World with and without China
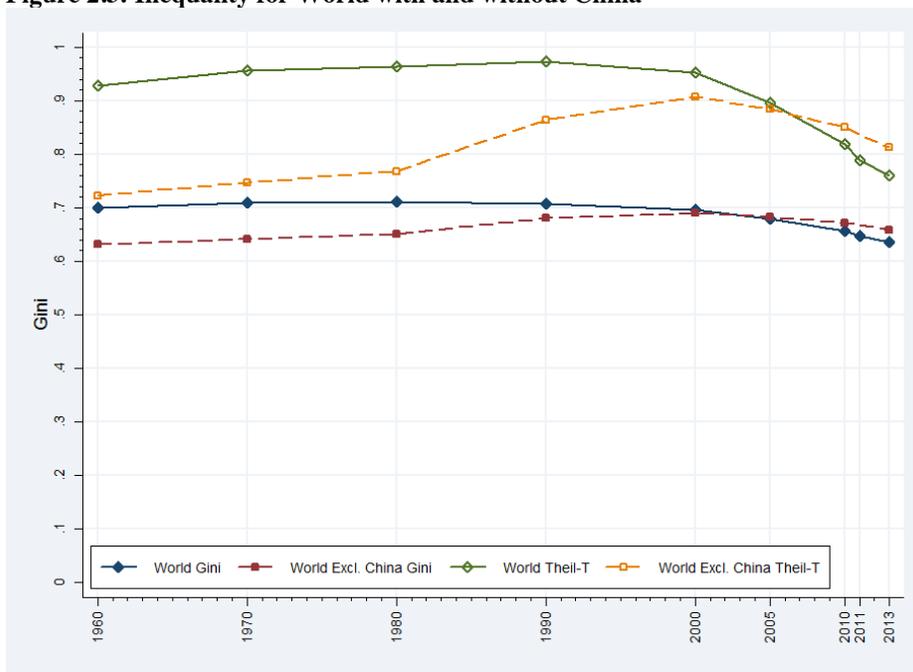
Source: Authors' calculations using GCD

2.6.2. Poverty Headcount Ratios

44

Figure 2.4 presents the evolution of the number of poor across the world (with and without China) defined by various poverty lines. Here again we see the huge impact of China in global poverty reduction. The rest of the world has seen the number of poor increase (or almost no change for the World Bank's $1.25 (2005 PPP) poverty line in recent years) from 1960 to 2010, with a slight decline after 2010, while China has seen a huge decline in number of poor. Future decline in poverty for this poverty line will need to be brought about as a result of a sizable contribution from other developing countries as the estimates poverty headcount ratio in 2013 for China for this poverty line is only eleven percent. Using GCIP, Reddy and Lahoti (2015) present alternate poverty estimates using 2011 PPPs and alternate poverty lines and compare them with the World Bank's latest estimates.

Figure 2.4 Poverty rates using various poverty lines
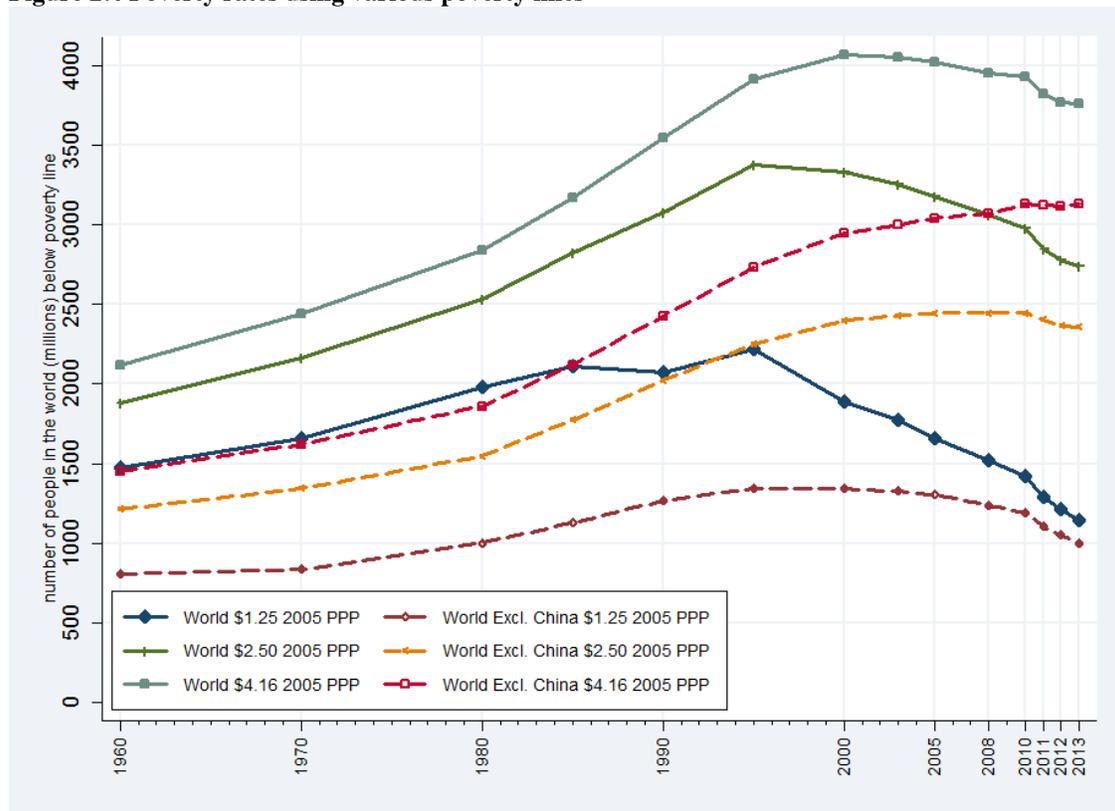
Source: Authors' calculations using GCIP.
Note: The $4.16 poverty line corresponds to the minimum per-capita expenditure required to obtain a home-cooked nutritious meal in the US as estimated in the Thrifty Food Plan of the United States Department of Agriculture (USDA). More details on the poverty line can be found in Reddy and Lahoti (2015).

These are only examples of the range of analyses to which the GCIP can be put. As we discuss in Jayadev, Lahoti and Reddy (2015a) the portrayal of how successful the development process

worldwide has been over the last fifty years depends very much on the criteria that we adopt (e.g. how high we set a poverty line). Such an insight is only possible because of the long-range and global analysis that the GCIP makes possible. Applications to monitoring and forecasting are also potentially interesting, since with continually updated and forecasted growth data, income and consumption profiles of national, regional and world populations can be estimated, so as to assess the impact of recent events or to examine the possible future consequences of trends.

2.7. Conclusion

The lottery of birth -- to whom one is born, when and where -- accounts for the majority of variation in the resources and opportunities available to human beings. Within nations, other influences -- one's gender, ethnic or racial category and other such factors -- serve to disadvantage some individuals in myriad, often invisible, ways from before they are born until their deaths. These patterns of inequality can be reinforced over generations through the effects of structural barriers, differences in political power or social discrimination, limiting the potential of persons to flourish equally or indeed to do so at all. One recent estimate suggests that the richest eight percent of individuals in the world enjoy the same income as the other ninety two percent of the population (Milanovic 2013) and this is likely an underestimate as the incomes of the rich are poorly reflected in household surveys and even in tax records[41].

In recent years, prominent social movements across the world (from the Indignados in Spain to the Movimento dos Trabalhadores Rurais Sem Terra (MST) in Brazil to the Occupy movements and the Arab Spring protests) have all been at least partly driven by the perceived illegitimacy of existing (and often growing) economic and political inequalities, and these are only the most well-known such instances. Governments in many parts of the world, it seems, are faced with dissatisfied or disaffected citizenries. At the same time, the middle class is also burgeoning in many countries and, especially if modestly defined, arguably in the world as a whole. Poverty appears to have fallen by certain measures although in a very geographically uneven way. Rapid economic growth appears to have been very unevenly shared within and across countries with some countries doing far better than others at generating inclusive growth, regardless of how this is understood. These diverse facts give rise to a complex picture of a changing global reality. Better research and data is needed to begin to capture the gross contrasts as well as the necessary nuances. Such data must be used not only for purposes of description but in order to better

---

[41] Some recent attempts have been made to try and include additional data from alternative sources such as tax and administrative records generally (most notable in this regard is the work of the authors of the Top Incomes Database). We hope to supplement future versions of our database with information from such sources and have begun initial work in this regard.

understand and influence the determinants of the changing relative and absolute fortunes of people.

We have described the Global Consumption and Income Project and presented some examples of the sorts of inferences that might be drawn from it. There are myriad applications that can be imagined. We present a work in progress that offers possibilities for a deeper understanding of the evolution of material well-being both within and across countries, for geographical regions and groups of countries and for the world as a whole, and that extends from description to explanation. It is to this end that we introduce our project, and invite interested specialists and the world public to help us build upon and improve it.

2. Appendix: Country Splits and Unifications

I. Countries undergoing Splits:

There are countries in our database which experienced splits and for which we have data pertaining to the entities both before and after this event. Examples of the affected pre-existing entities include Czechoslovakia, Ethiopia, Indonesia, Malaysia, Pakistan, the USSR and Yugoslavia. Other countries which experienced splits over the period, such as Sudan, are ones for which we do not yet have sufficient data to incorporate them separately.

I.1 Former Socialist Economies:

For Czechoslovakia, the USSR and Yugoslavia we apply the following procedure. In most of these cases, we have data on means and distributions for the countries emerging from the split from near to the year in which they were formed and some distributional data on the combined country for the earlier period. As some of the countries undergoing splits were also formerly socialist economies, problems of available of distributional and mean information are compounded by the often lesser availability of household data for such countries.

Our approach to handling these countries is to attempt to assign a pseudo-mean and pseudo-distribution to the constituent countries even prior to the split, while recognizing that the assigned values may in fact be more characteristic of the unified country, and thus recommending appropriate interpretative caution. Such an approach allows us to estimate trends for the individual constituent units over a longer period, as well as to construct and report aggregates (whether prior to or after the split). Each of these possible choices has its benefits and costs in terms of statistical meaningfulness vs. inter-temporal comparability.

For the affected countries mentioned above, we use distributional data from the combined country to create a pseudo-distribution of each of the countries undergoing splits for years when the countries were one, as we do not possess distributional data for the individual constituent countries. We recognize that this is an inadequate assumption, in part because the distribution for the unified country reflects differences in income between successor countries as well as within them.

Although we do possess some mean estimates for the combined countries from national accounts or independent academic and institutional estimates, we do not use these for the constituent countries both because this would mean using the same mean for all constituent countries in the earlier years and because we don't have reliable national-level inflation data for the pre-split period to convert these estimates into units of a common PPP base year after the split. We therefore instead estimate means for the countries undergoing splits by extrapolating backward using real per-capita-income growth rate data, which is available in many instances. Where growth rate information for the constituent countries is not available, we use the growth rate of the combined country as a proxy for the growth rate of each of the countries resulting from the split. In particular, for the three countries undergoing splits mentioned

Table 2.A1: List of countries that split between 1960 and 2015

	Survey Years	Growth Data	Source of Growth rate data	Inflation Data
Yugoslavia	1968-1990	1961-2006	Angus Maddison	
Bosnia and Herzegovina	2001-2007	1991-2015	TED and WDI	1994-
Croatia	1988-2008	1981-2015	TED and WDI	1985-
Montenegro	2005-2011	1991-2013	TED and WDI	2000-
Serbia	2002-2010	1991-2013	TED and WDI	1994-
Slovenia	1987-2011	1981-2015	TED and WDI	1992-
Macedonia, FYR	1994-2008	2015	TED and WDI	1990-
Kosovo	2003-2006			2000-
USSR	1980-1989	1961-1990	Angus Maddison	
Armenia	1996-2012	1981-2015	TED and WDI	1990-
Azerbaijan	1995-2008	1981-2015	TED and WDI	1990-
Belarus	1988-2011	1981-2015	TED and WDI	1990-
Estonia	1988-2011	1981-2015	TED and WDI	1992-
Georgia	1996-2012	1966-2015	TED and WDI	1965-
Kazakhstan	1988-2010	1981-2015	TED and WDI	1990-
Kyrgyz Republic	1988-2011	1981-2015	TED and WDI	1987-
Latvia	1988-2011	1966-2015	TED and WDI	1965-
Lithuania	1988-2011	1981-2015	TED and WDI	1990-
Moldova	1988-2011	1981-2015	TED and WDI	1989-
Russia	1988-2010	1961-2015	TED and WDI	1989-
Tajikistan	1999-2009	1981-2015	TED and WDI	1985-
Turkmenistan	1988-1998	1981-2015	TED and WDI	1987-
Ukraine	1985-2010	1981-2015	TED and WDI	1987-
Uzbekistan	1988-2003	1981-2015	TED and WDI	1987-

Czechoslovakia	1964-1992	1961-2006	Angus Maddison	1975-
Czech Republic	1988-2011	1971-2015	TED and WDI	1990-
Slovakia	1988-2011	1986-2015	TED and WDI	1992-

above, we employ data on growth rates for the constituent countries from 1980 onwards from the Total Economy Database. Prior to 1980 we use the growth rates for the combined entities as provided by Angus Maddison.[42]

Table A1 has details on the data used for each of the combined countries undergoing splits.

I.1 Other Countries Undergoing Splits

For the other countries undergoing splits, we have made specific assumptions, for instance about the coverage of surveys from before and after the split. These are mentioned in the online appendix on country assumptions (see gcip.info).

II. Countries Undergoing Unification:

Among the countries undergoing unification during the time interval covered by the database are Germany and Yemen. We do not have sufficient information for the constituent parts of Yemen prior to or posterior to its unification to form a picture of the country at this time.

For Germany, we use West Germany's distribution and mean for all of Germany prior to unification. We are actively interested in finding and integrating East German data from prior to unification so as to improve upon this inadequate approach.

---

[42] Maddison provides estimates of the level of real per-capita income in 1973 and in 1990 for a number of constituent countries (e.g. Soviet Republics) which could permit determining an average annual growth rate over the period, but we refrain from employing a smoothed estimate in favor of using TED annual estimates for the period from 1980 to 1990.

3. Who Got What, Then and Now? A Fifty Years Overview from the Global Consumption and Income Project [1]

Abstract

Using new comprehensive data and tools from the Global Consumption and Income Project (GCIP), covering most of the world and more than five decades, we present a portrait of the changing global distribution of consumption and income and discuss its implications for understanding inequality and inclusivity of growth and development. We show how regional distributions of income and consumption have evolved very differently over time. We undertake sensitivity analysis to quantify the impact of various choices made in database construction and analysis. We find, in keeping with other studies that levels of consumption and income have increased across the distribution, that the global distribution has become more relatively equal due to falling inter-country relative inequality though more absolutely unequal. We describe the singular role of China in bringing about the observed changes. We also present new findings about some old 'stylized facts' as well as new ones deserving study.

3.1. Introduction

The distribution of income within and across countries has been the subject of considerable and growing academic and popular interest in the recent past. This has been driven both by moral concerns and by the need to understand the effects of global economic integration in order to make political and institutional choices. Debates about the evolution of poverty, inequality and the extent to which the benefits of growth and development have been shared both within and across countries have all drawn attention.

The last three decades have also experienced a data revolution. More and better surveys have been collected by multiple agencies, including the World Bank (e.g. through the Living Standards Measurement Surveys), the LIS (formerly known as the Luxembourg Income Study) and others, with some effort being made to achieve greater international comparability. Extensive efforts to collect price data by the International Comparison Program (ICP) have also

---

[1] This paper is joint work with Arjun Jayadev and Sanjay Reddy.

provided greater ability to account for differences in purchasing power across countries (if not over time). The ICP collects comparative price and expenditure data from countries to estimate purchasing power parities (PPPs) of countries periodically. PPPs are meant, *inter alia*, to facilitate cross-country comparisons of material well-being through better assessment of differences in price levels and resulting command over real resources (Deaton & Heston, 2010). As a result of these developments, and accessible computing power, researchers have come to be able to produce estimates of global inequality based on collections of household surveys -- as opposed to on estimates of GDP per capita[2] -- in the last twenty years. There have been in recent years additional attempts (see e.g. Milanovic, 2012; Lakner & Milanovic 2015; Anand & Segal 2014; Edward & Sumner 2015 for useful examples of such work) to provide more developed portraits of the pattern of material living standards of the world's population.

Our paper adds to the literature both in terms of methods and substantive conclusions, building on a new and unprecedented resource —the Global Consumption and Income Project (GCIP) - that is composed of two distinct databases - the Global Consumption Database and the Global Income Database. These databases improve upon previous efforts in a few important ways. Most existing studies mix income and consumption surveys and use this pooled data to compare inequality across countries or regions, and over time. By contrast, to avoid the potentially misleading effects of such pooling, we standardize the welfare concepts used. In particular, to portray the global consumption distribution we use only distributions from consumption surveys or estimated 'equivalents' implied by adjusting income surveys that are part of the larger database. We apply an analogous principle when portraying the global income distribution. Another limitation of current databases is that they often provide estimates for only a few benchmark years or restrict coverage to a subset of countries. Although there could be reasons for such an approach, the GCIP in contrasts attempts to maximize coverage (by the judicious use of interpolation and by using all primary data sources available to us). This makes it possible to provide a considerably more comprehensive portrait by covering a larger group of countries (161), a longer time period (1960 to the present as opposed to the mid- or late 1980s to the present) and annual portraits of the world distribution, drawing on the nearest available surveys. In addition, in the GCIP, we employ purpose-built tools for aggregation that allow us to analyze evolution of material well-being for sets of individuals belonging to user-defined country groups (e.g. regions of the world or countries sharing other specified traits) in a manner that is not possible with other existing databases. These tools require the annual estimates we generate, as without them the resulting holes in the data would make aggregation impossible. We extensively

---

[2] Early estimates of global differences in income based on per-capita income data go back at least to Nurkse (1953) who draws on early data from the United Nations.

test the sensitivity of our results to different assumptions, including about regarding the selection of the data for the consumption and income databases, standardization method, exchange rates used and the impact of exclusion of certain countries from the global distribution. The approach we adopt is described in detail in Lahoti, Jayadev and Reddy (2016) and in the complementary materials available on the associated website (gcip.info).

In this paper, we undertake two separate but interlinked tasks using the dataset and the methods and tools we have developed along with it. First, we provide a portrait of the global distribution of income [in what follows we will often refer to income as shorthand for 'income and consumption'] and its constituent regional distributions and describe their evolution over time. Second, we describe the degree to which income or consumption levels of individuals have increased at various points in the distribution. In particular countries, most notably China, beneficial effects of rapid growth such as poverty reduction appear to have been combined with very inequitable distribution of gains. We undertake a special focus on the countries that are experiencing rapid growth to identify whether this is the case more generally, and to characterize their contribution to the overall change in the global distribution.

To briefly summarize our key results, a few points may be noted. First, global incomes have increased across the board over the last three decades. Median income doubled from 2005 PPP $73 to 2005 PPP $154 per month from 2000 to 2015 after much slower growth in the previous decade. Second, the global income distribution has become substantially more relatively equal in the last three decades. Most of these improvements came in the period after 2000 and can be attributed to inter-country rather than intra-country effects: the very rapid growth of China and to a lesser extent of India during this phase. China's population has experienced an enormous advance that has shaped the world distribution as a whole and others' relative positions. The rest of the world excluding China has witnessed increase in median incomes of only twenty six percent between 1980 and 2013 while China's median has increased by 1188 percent in the same period; and the income distribution for the rest of the world is still twin peaked. Third, country experiences vary widely, and there is some evidence of 'inequality convergence' with previously more unequal countries becoming more equal over time and vice versa, although the majority of the world's population lived in countries with rising inequality. Also countries from Sub-Saharan Africa have the highest income inequality in the world upending the `stylized fact' that Latin American countries are among the most unequal in the world[3]. Fourth, we focus on the fifteen

---

[3] The main reason for this somewhat unusual finding is that the current 'stylized' fact derives from the artifact that Sub-Saharan countries typically are based on consumption surveys (which are more equal than income surveys) while the opposite is true of Latin American countries. Correcting for this leads to the conclusion that Sub Saharan Africa is both the poorest and most unequal major region of the world.

fastest growing countries in the world and find that despite significantly high growth across all these countries, the impact of Chinese growth dominates in changing the world income distribution. Finally, we undertake a welfare analysis and find that while global 'welfare' estimated by any standard income-focused aggregative welfare function has increased substantially over the last fifty years most of the change has happened in the period since 2000. The rest of the world excluding China has seen no Pareto-improvements in the period 1970 to 2000.

3.2. Data

We use data from the Global Consumption and Income Project (GCIP), constructed according to the methods presented in Lahoti, Jayadev and Reddy (2016) and associated online materials (available on gcip.info) which together offer details of the procedures used, ongoing revisions to the methods employed, and how the data differs from other available sources. We construct estimates of annual data from 1960 to 2015, for each percentile of the population from 161 countries covering 97 percent of the global population in 2015. The GCIP is a complete 'time-space system' which produces estimates for every country-year, which is essential in order for us to be able to use it in a flexible way to construct estimates for country aggregates.

The GCIP was also, distinctively, built keeping the goals of attaining transparency, replicability and flexibility foremost during its construction. We aim to document fully the assumptions and choices made in the database generation process. The database is constructed in a manner that is intended to make it possible to adopt alternate assumptions and thus to test the sensitivity of the choices made, of which we demonstrate some examples in this paper.

Construction of the GCIP datasets involves several decisions on selection of data and methods used; some of the most important ones are discussed here briefly. First, we restrict ourselves to surveys that provide household per-capita data, as data employing equivalence scales in their construction use widely variable and incomparable methods and constitute a smaller proportion of the available data[4]. Second, for country-years with no consumption and income survey we

---

[4] The majority of developing countries report grouped survey data in per-capita terms and hence we cannot apply any other equivalence scale to such grouped data. In the case of developed countries where data might be reported using an equivalence scale (for example for surveys from LIS, SILC and ECHP) we have used the underlying unit data to estimate instead per-capita measures. Findings based on applying an equivalence scale might differ from the ones based on using per-capita data as distributions of size and composition of households can be systematically different across regions. For example, Sub-Saharan African households tend to be larger and with more children than those in Latin America or other regions. But there is no consensus on the appropriate equivalence scale to use and how this would vary by country or region. Determining an appropriate equivalence scale is outside the scope of this paper, although as we note in Jayadev, Lahoti and Reddy (2016) the choice of scale is greatly consequential and the arguments for choosing one of these are weak. In order to maintain comparability we use per-capita measures as is also done in Lakner and Milanovic (2015) and in the World Bank's Povcalnet database.

interpolate or extrapolate the consumption or income profile using survey data from the closest survey years and appropriate growth rates[5] from the national accounts. This method is similar to

Table 3.1: Summary Statistics for Surveys in Global Consumption Database (GCD)

	1960-69	1970-79	1980-89	1990-99	2000-09	2010-13	Total
Number of Surveys	79	81	283	518	700	337	1998
Number of Countries with Surveys	44	44	98	132	152	119	161
% Consumption Surveys	14	12	23	39	45	39	37
% Surveys Covering Complete Population	56	68	77	98	100	100	93
% Surveys Covering all Areas in the Country	95	95	93	96	97	98	96
% Surveys with Mean Levels Information	41	49	65	92	96	92	86
Source of Surveys (%)							
ECHP	0	0	0	17	2	0	5
EU-SILC	0	0	0	0	25	34	14
LIS	3	15	13	8	5	7	8
Povcalnet	0	0	15	31	40	35	30
Data from Statistical Office	0	0	0	0	1	2	1
SEDLAC	0	1	7	17	21	19	16
WYD	0	0	1	2	1	0	1
WIID	97	84	64	26	6	2	26
Regions (%)							
East Asia & Pacific	5	14	11	9	10	11	10
Europe & Central Asia	38	43	54	46	51	52	50
Latin America & Caribbean	23	20	19	23	21	22	22
Middle East & North Africa	5	6	3	3	3	2	3
North America	0	5	1	1	1	1	1
South Asia	19	7	4	3	3	3	4
Sub-Saharan Africa	10	5	8	13	11	10	11
Income Grouping (%)							
Low income	9	5	4	11	19	13	9
Lower middle income	31	17	24	29	15	7	24
Upper middle income	39	45	44	33	42	53	40
High income	21	33	28	27	25	27	27

---

[5] We use growth rate of household final consumption expenditure per capita when available for interpolation/extrapolation. We do not use any adjustment factor for the growth rate as used by Ferreira et. al (2016) but this should have minimum impact on our interpolations as means are bounded from both sides.

the one used in Povcalnet and is described in detail in Lahoti, Jayadev and Reddy (2016)[6]. These extensions might raise some valid concerns that we are aware of and hence are careful in our assumptions and provide an option for concerned users to use only actual survey data. For the period after mid-1980's the density of surveys is dense, and our interpolation/extrapolation is unlikely to impact results in any significant way. A larger amount of the data before 1980 is interpolated or extrapolated due to sparse survey information and thus has to be treated with caution, but is still indicative of the trends during this period.

Third, the GCIP uses a regression-based 'standardization' method to predict the consumption shares of each quintile of the population for the Global Consumption Database (GCD) in country-years, which have an income survey but no consumption survey and the obverse for the Global Income Database (GID). The distribution of income is known to be more concentrated than that of consumption, but almost none of the existing databases correct for this (Ferreira et al, 2016; Lakner & Milanovic, 2015). We use information from 204 country-years across 71 countries in which there is both consumption and an income survey reported by the same statistical agency to derive a regression relationship between consumption and income for each quintile, using a Seemingly Unrelated Regressions (SUR) approach. We use various controls and different methods and specifications to test robustness of the relationship between income and consumption for each quintile. The set of countries on which we have information are spread across geographical and income groupings and over time. We find that the regressions are fairly accurate for in-sample prediction. We thus apply the same relationship to make out-of-sample predictions for other country-years. After doing so we find that the difference between income and consumption Gini coefficients (arising from comparisons of results from the two standardized datasets for the same country-year) is between 7 and 10 points, which is comparable to estimates in the literature based on simply comparing average differences[7]. A detailed description of the method used for standardization is provided in Lahoti, Jayadev and Reddy (2016)[8].

---

[6] Almost all databases, even the ones using benchmark years do some degree of interpolation/extrapolation to line up surveys for a particular year to do comparisons across countries.

[7] In the full PovcalNet data, the average Gini index of consumption surveys is approximately 10 Gini points lower than the average Gini over the income surveys (Lakner and Milanovic, 2015). Li, Squire, and Zou (1998) suggest an adjustment of 6.6 Gini points based on comparison of average values in the Deininger and Squire (1996) database.

[8] We do not standardize the various income concepts (gross, disposable, monetary) as we do not have enough data to ascertain relationships between gross vs. disposable vs. monetary shares for different quintiles as was the case with consumption and income. When we have a choice between surveys with various income concepts for a country-year, we prefer the disposable income over other income concepts (see Lahoti, Jayadev and Reddy (2016) for more details on our preference ordering).

Figure 3.1: Survey Means and GDP Per Capita for Survey Years from 1960 to 2013
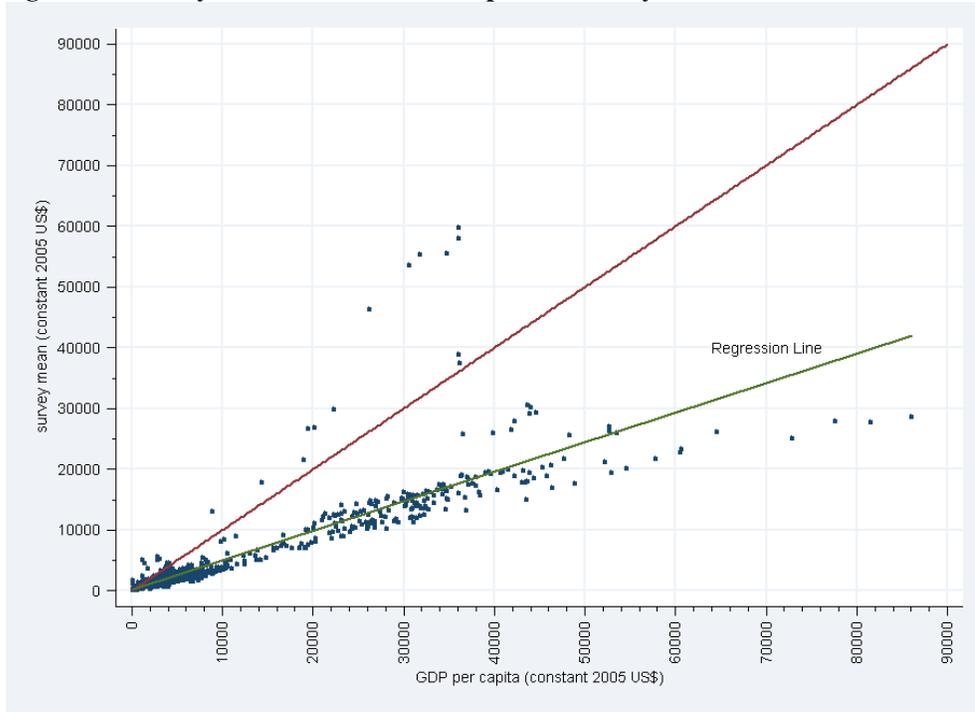
Finally, we use estimates of consumption or income levels from surveys wherever they are available. This is a consequential choice, since survey means are often discrepant from (and typically lower than national accounts means). Figure 1 shows the relation between GDP per capita and survey means for country-years in our dataset. It can be seen that, generally, survey-means are much lower than the household consumption component of GDP-per-capita and thus of GDP-per-capita itself[9]. Investigation of the data suggests that this is a phenomenon that crosses decades and world regions. For this reason among others, the estimates of the absolute level of income that we arrive at, as well as its distribution, must be viewed with the proverbial salt in hand. We also standardize the means by converting means from income surveys (whose distribution was standardized in the previous step) used in the Global Consumption Database (GCD) to consumption means by using the share of consumption in National Accounts as a multiplicative conversion factor, while we use the reciprocal share to convert consumption to income means for the Global Income Database (GID).

---

[9] The growth rates of survey means and of GDP-per-capita are also substantially different (but it is more difficult to find a clear trend of growth in survey means vs. in GDP-per-capita, as shown in our discussion of fast growing countries later).

We present the main results using 2005 PPPs and not the most recent 2011 PPPs as there is still debate in literature on which ones are more appropriate (Deaton and Aton, 2017; Ravallion, 2014) and we ourselves take an agnostic view with respect to PPP base years, especially where comparison over lengthy periods of time as well as space are involved. However, we have reproduced all of our primary results using both PPP base years and report the differences where they are noteworthy.

The Global Consumption Database (GCD) has a wide geographical and temporal coverage with data from various secondary and primary sources as seen in Table 3.1. The GCIP presently contains survey data for 1998 country-years spanning the period of 1960-2013 for 161 countries. Most of the surveys are nationally representative (96% cover complete geographical area and 93% cover the entire population). Our data is drawn from various secondary sources like the World Bank's Povcalnet database, UNU WIDER's World Income Inequality Database (WIID), Socio-Economic Database for Latin America and Caribbean (SEDLAC), European Community Household Panel (ECHP), Statistics on Income and Living Conditions (EU-SILC), World Income Database (WYD) and from the LIS database and from data directly obtained from National Statistical Offices.

The density of surveys varies drastically across the decades. The 1960's and 1970's have the lowest density of surveys with only 79 and 82 country-year observations respectively from 44 countries. This is largely because of paucity of household surveys, especially in the developing countries, during this period. Our choice of using only per-capita surveys also restricts the number of country-years as this period has several surveys where only total income is reported at household level with no adjustment for household size. The formerly communist countries also have sparse data on income or consumption distribution prior to 1990. Given this data limitation we advise caution in interpreting the results encompassing the earlier period.

The GCIP datasets used in this paper reflect further improvements in methods of construction of the dataset described in Lahoti, Jayadev and Reddy (2016). In particular, we have made systematic efforts to reduce the volatility of data arising from surveys being selected from different data sources, by using specific methods to select surveys that belong to consistent long survey series. We have also extended the datasets by adding new survey sources (European surveys from SILC and ECHP) and incorporating surveys from recent years beyond 2010.

3.3. Inequality

We begin by presenting a portrait of the distribution of global consumption and income using the GCID. All studies on global inequality combine data across different countries that vary in terms of the welfare concept as well as in what is included operationally in the welfare concept[10]. In particular, items that are included in consumption/income surveys are not consistent across countries or over time. One important example, which by itself could very substantially influence comparisons, is imputed rent for self-owned homes, that is included in some countries but not all and even when included is estimated based on different methods across countries. Such variations can be sizable across different survey sources even for a single country, leading to difficulty in making comparisons over time, or to apparent volatility due to the combination of

Table 3.2: Global Consumption and Income Levels

Consumption per capita per month at various percentiles in 2005 PPP							
year	10th	20th	50th	90th	95th	99th	Mean
1960	6	10	36	324	496	893	122
1970	7	11	42	453	686	1226	163
1980	10	15	44	530	795	1395	189
1990	16	22	50	603	933	1678	215
2000	19	27	61	610	955	1800	231
2005	23	32	74	639	1005	1931	252
2008	24	34	82	661	1028	1969	269
2010	26	37	90	692	1069	2013	281
2013	31	44	105	719	1070	2014	295
2015	34	48	116	759	1117	2091	313
Income per capita per month at various percentiles in 2005 PPP							
year	10th	20th	50th	90th	95th	99th	Mean
1960	8	12	40	408	653	1219	155
1970	9	14	44	570	904	1662	207
1980	12	18	47	681	1052	1907	242
1990	14	22	60	765	1176	2173	270
2000	16	26	73	830	1290	2414	297
2005	20	33	95	894	1378	2577	333
2008	22	36	111	947	1434	2683	363
2010	24	39	121	981	1474	2749	378
2013	28	45	139	1016	1517	2837	402
2015	31	50	154	1080	1599	2981	430

Source: Authors' calculations.
Note: All numbers are in 2005 PPP $.

---

[10] Even though an expert group has recommended guidelines for collection of income distribution data (de Canberra, 2001) they are rarely adhered to in developing countries.

such sources. Global estimates are also prone to several sources of uncertainty both intrinsic to the methods used and arising from the plausible use of different methods. Areas in which such uncertainties arise include survey design and sampling, PPP estimation, standardization methods used to compare income and consumption and others. Although they are known to exist, not all of the resulting probable errors can be quantified. Given these difficulties, all estimates based on the existing data, both ours and others, are subject to concerns about meaningfulness and reliability (see Anand & Segal, 2008, 2014 for more details). It is interesting that nevertheless, certain regularities emerge across different research efforts.

Figure 3.2: Global Income Distribution
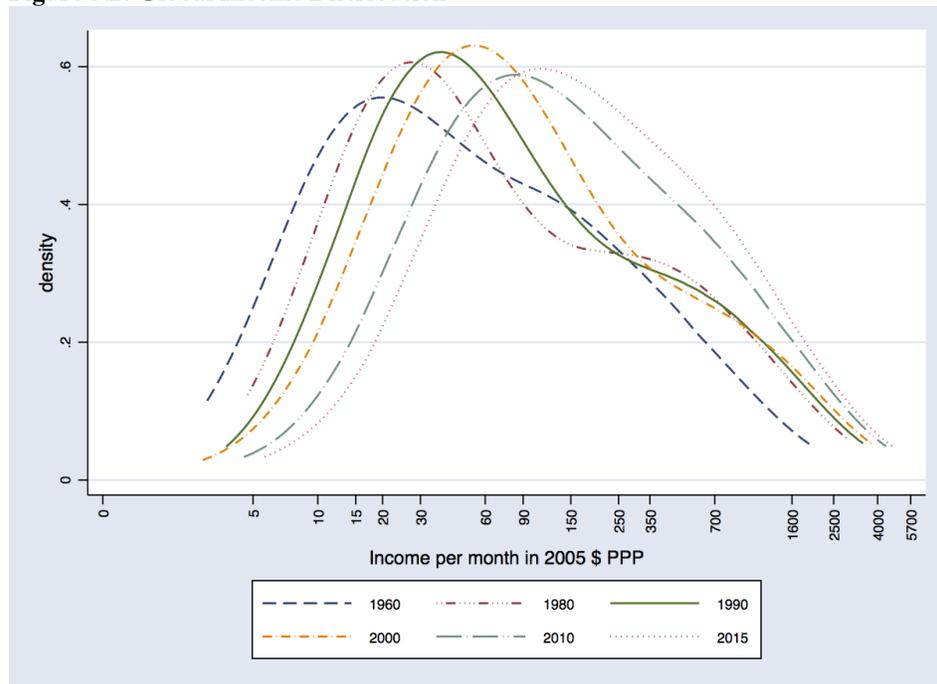
Note: kernel=epanechnikov, stata default bin size.

For example, our database is able to reproduce the general finding (see e.g. Milanovic, 2012; Lakner & Milanovic, 2015) that the last three decades have seen substantial increases in material living standards for the world population considered as a whole, especially in the period after 2000, when comparing percentiles with corresponding percentiles. In 1990, half of the world's population had consumption levels of $50 (in 2005 PPP dollars) per month or $1.64 per day, or less. By 2000, this fraction had fallen to 43% and by 2013 fell further to 24%.

Table 3.2 provides snapshots of the distribution of global consumption and income distributions for different years between 1960 and 2015[11]. The extraordinary growth rates experienced by a number of developing countries in the 2000s are evident in the fact that while the median income grew at a compound annualized rate of 2 percent from 1990 to 2000, it more than doubled between 2000 and 2013 to 4.2 percent.

Figures 3.2 shows the evolution of the distribution of global income over a longer period of time. We find the gradual disappearance of the 'twin-peaks' phenomenon of slight bi-modality in the global distribution studied by, among others, Milanovic, 2005; Weisbrod et al, 2007; Lakner & Milanovic 2015. It is crucial to note that in this figure and similar ones that follow the scale is a logarithmic one in which given visual distances on the scale correspond to larger and larger differences in absolute income as one examines comparisons involving higher incomes and may therefore visually understate the level of global inequality.

The remarkable impact of Asia's rise in the global economy over the last 25 years can be seen by examining the regional superposition graphs (Figures 3.3 and 3.4) for consumption in 1990 and in 2013. These show the contribution of each geographical region to the overall distribution of consumption as well as their individual distributions at two periods in time. The most striking development is perhaps the disappearance of pre-existing twin peaks in East Asia and the movement of most of its population to a higher level of consumption.   As the figures show, the median of the distribution in East Asia was at $32 2005 PPP per month in 1990 but by 2013, the median of the distribution had come to be around $135 2005 PPP per month. As is also clear, despite India's substantial GDP-per-capita growth in the period, compared to 1990, South Asia has fallen behind East Asia and the Pacific[12]. The median of the distributions was roughly the same in 1990, but by 2013 the median in South Asia is at less than half the level of East Asia and Pacific. Another notable feature of the data is a degree of convergence within Europe and Central Asia leading to a less skewed and more even distribution, perhaps due to increases in income in formerly planned economies.

As a result of the fast growth of the Chinese (and to a lesser extent Indian) economy, the world distribution of income and consumption is more equal in 2013 than it was in 1990. Table 3.3 provides estimates for global interpersonal income inequalities over the period for various

---

[11] Note that the data before 1980 should be treated with extreme caution, since there are only survey-based observations for around a fourth of the countries in the 1960s and 1970s.

[12] This is partly because growth rate of survey means in India is far less than its growth in GDP-per-capita (Annualized compound growth rate of 2.93 percent in survey means vs. 4.67 percent in GDP-per-capita from 1990 to 2013).

Figure 3.3: Regional Superposition Consumption Density for 1990
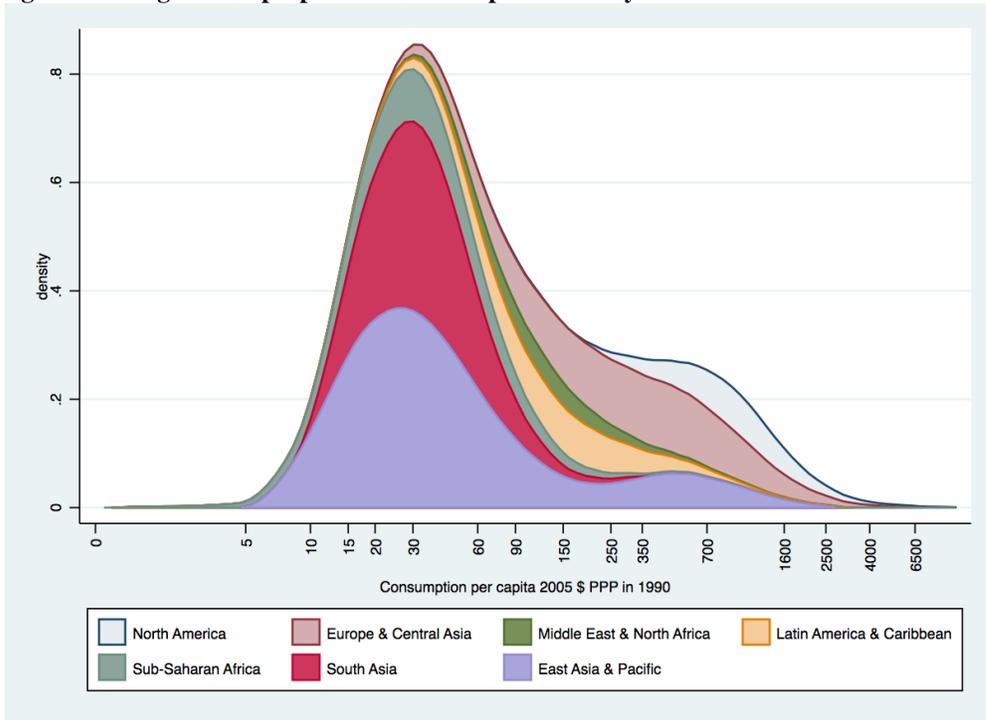
Figure 3.4: Regional Superposition Consumption Density for 2013
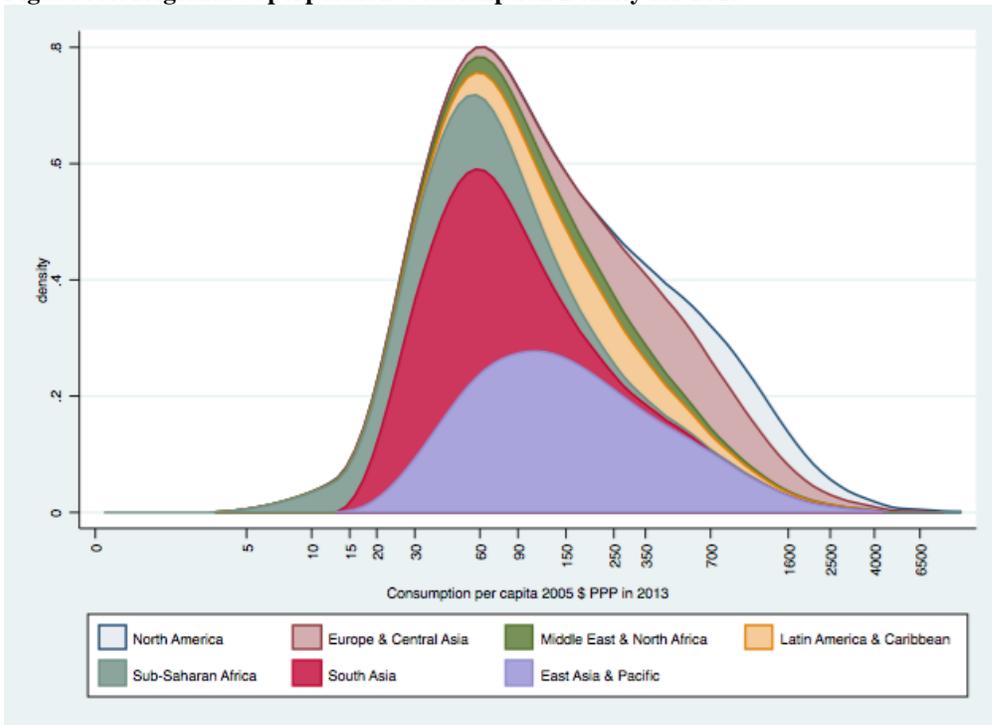
inequality measures. In 1990 the global income Gini coefficient was 0.72 while the global Theil index was 1.02. In 2013, these had fallen to 0.66 and 0.82 respectively. Even if one were to exclude China, global inequality would be estimated to have fallen although only due to the interval between 2005 and 2013. Once again it is noteworthy that all the reduction in inequality occurred after the 2000s. In the period 1990 to 2000 global relative inequality did not decrease substantially.

Table 3.3: Global Relative Income Inequality Measures

	Gini	MLD	Theil-T	Palma Ratio	Mean to Median Ratio
1960	0.71	1.12	0.97	16.02	3.89
1970	0.72	1.22	1.01	19.12	4.71
1980	0.72	1.2	1.01	17.84	5.13
1990	0.72	1.17	1.02	16.37	4.54
2000	0.72	1.13	1.01	14.99	4.05
2005	0.69	1.04	0.94	12.75	3.51
2010	0.67	0.98	0.86	11.34	3.12
2013	0.66	0.93	0.82	10.19	2.9

Source: Authors' calculations.

Figure 3.5: Change in within-country inequality in the period 1980-2013
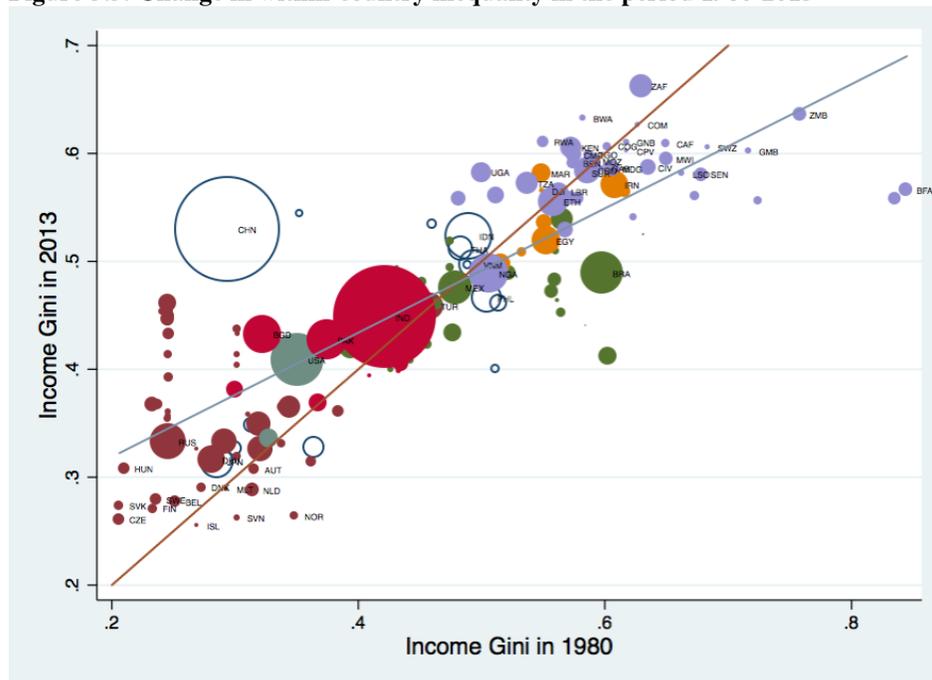
While the world distribution of income has grown more equal, the experience of within-country inequality has been more disparate. Figure 3.5 provides evidence of the change in the income

Gini coefficient within country between 1980 and 2013. The size of each circle corresponds to the population of the country in question. Where observations are above the 45 degree line, countries have become more unequal, while countries have become less unequal if the observation lies below the line. There is some evidence of inequality convergence noted by Ravallion (2001), Dhongde and Xing (2013) among others as the line of best fit for the observations is flatter than the 45-degree line   Roughly half of the countries show rising inequality, but about two-thirds of the population lives in countries with rising inequality (a difference mainly driven by China).

One interesting point to note about the scatterplot on inequality convergence is the preponderance of Sub-Saharan African countries among those with high inequality both at the beginning and at the end of the period (higher than most Latin American countries, contrary to a prominent received 'fact' in economics). This is a result of the fact that we have standardized the concept of advantage (consumption vs. income) so as to make estimates of inequality more comparable; while surveys in Sub-Saharan Africa typically are of consumption, those in Latin America are typically of income, which without adjustment leads to greater estimated inequality in the latter.   The high-inequality of Sub-Saharan African countries (and the finding that it is among the highest in the world, and in particular higher than in Latin America) holds in our database for earlier decades as well. In addition, it is interesting to note the extent to which previously relatively equal European economies have become more unequal.

Table 3.4: Average Within Country Income Inequality by Region

	Unweighted			Population weighted		
	1990	2000	2013	1990	2000	2013
East Asia & Pacific	0.45	0.47	0.46	0.47	0.5	0.5
Europe & Central Asia	0.3	0.34	0.36	0.32	0.35	0.36
Latin America & Caribbean	0.49	0.51	0.48	0.53	0.54	0.48
Middle East & North Africa	0.53	0.54	0.53	0.55	0.55	0.54
North America	0.35	0.37	0.37	0.37	0.4	0.4
South Asia	0.41	0.44	0.41	0.44	0.44	0.44
Sub-Saharan Africa	0.62	0.6	0.58	0.6	0.59	0.56
World	0.46	0.47	0.46	0.45	0.47	0.47

Source: Authors' calculations.

Although global inequality has declined between 1990 and 2013, within-country inequality increased on average in the same period (Table 3.4). The population-weighted average income Gini coefficient of countries increased. As we showed previously considerable re-ranking can

occur when one uses a consistent income Gini coefficient (obtained after standardization). We observe this in case of Latin American and Sub-Saharan countries that reverse positions as a result of this. Average within-country income Gini coefficients for Latin American countries are higher than consumption Gini coefficients for Sub-Saharan African countries, but Sub-Saharan African countries have higher income Gini coefficients after standardization than Latin American countries[13]. Sub-Saharan Africa and Latin American countries also witnessed substantial decline in average within-country inequality between 2000 and 2013, while other regions saw either an increase or no change in within-country inequality.

Figure 3.6 provides a depiction of inequality within countries as measured by the income share of the top decile. We should note here that in every country this is a likely under-estimate due to the well-known problem of under-sampling, under-reporting and higher non-response among the very rich. This noted, a number of Sub-Saharan African nations dominate the map.

Figure 3.7 shows the total world inequality attributable to between and within country inequality (using an additively decomposable income inequality measure). Again, as a result of the growth of China, the importance of between-country inequality has fallen sharply in explaining global inequality. Within-country inequality in 2013 accounts for 45% of global inequality as opposed

Figure 3.6: Income Share of the Top 10% in 2013 (based on surveys, 2005 PPP)
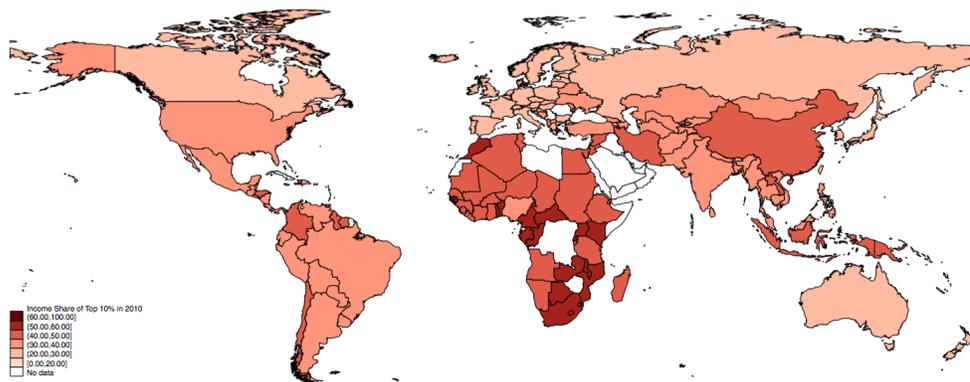
to 24% contribution it made in 1980. In fact, the between-country component of inequality is at its lowest ever in our sample. Some of this convergence is driven by the fast growth of China and other countries but it also has to do with the steady, even if relatively slower, growth of other developing countries (such as Brazil). The contribution of China in this regard is crucial.

---

[13] Sub-Saharan African countries consumption inequality is also higher than consumption inequality in Latin American countries.

Without China, the between-country component of global interpersonal inequality would have fallen marginally from 67 percent in 1990 to 62 percent in 2013 (as compared to falling more markedly from 76 percent to 55 percent with China).

Given the facts of global income inequality convergence and the fall in between country inequality over the last 50 years, it is perhaps not surprising to see, as we do in Figure 3.8 that global relative income inequality is at its lowest measured, whichever inequality measure we use. Even when one examines global absolute inequality (relative inequality as measured by any inequality measure such as the Gini or Theil times the mean), which gives consideration to absolute differences, the remarkable equalizing impact of Chinese growth in the recent past is visible. Figure 3.9 shows the trend in absolute inequality over the last 40 years. Unsurprisingly, given that these measures register increasing inequality as long as the absolute income increase going to a poorer person is lower than that going to a richer person, over most of the period it shows rising inequality (Atkinson & Brandolini, 2010; Niño-Zarazúa, Roope & Tarp, 2016).

Figure 3.7: Decomposition of Global Income Inequality into within and between country components (2005 PPP)
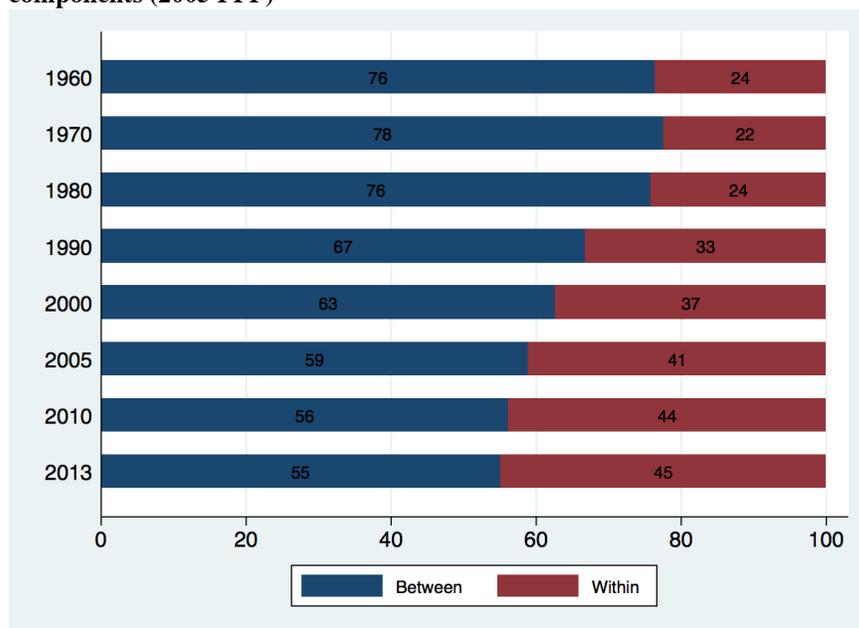
Note: Inequality measure is Mean Log Deviation - GE(0)

Despite the very high bar for inequality reduction, there has been a relative decline in increase in inequality during the last period of high overall developing country growth (2005 and 2010), if one uses the Theil measure or the Absolute Mean Log Deviation measure of inequality.

Of course, there are many complex and diverse experiences across regions in regard to growth and distribution in this period. In order to capture these variations, we undertake a regional decomposition of inequality trends (Figure 3.8). We use the World Bank's regional categories and assess the trends in income Gini coefficients for the resulting regions considered as a whole. A few striking conclusions emerge. First, the East Asia and Pacific has been one of the highest inequality regions in the world till 2005 (contrasting with the widely held 'stylized fact' that individual East Asian countries have had fairly low inequality as compared to other regions historically). This is because it contains populations at very different levels of development, from very poor Chinese or Indonesians for example, to Japanese (and more recently Koreans or Singaporeans) who enjoy levels of income of rich countries. However, even within this disparate group, China's rapid income growth has meant that inequality has fallen since 1980 because of the impact of that income-growth on inter-country differences within the region.

Figure 3.8: Regional Relative Interpersonal Income Inequality (Gini, 2005 PPP)
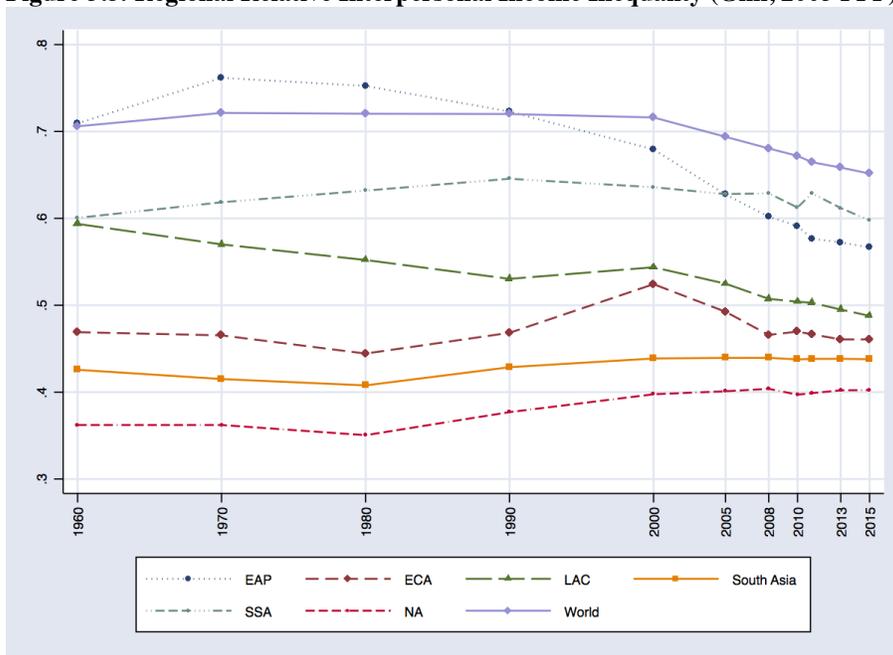
Meanwhile, at the other end of the spectrum, North America is much more homogenous as a region and overall inequality is correspondingly lower. However, rising inequality within the US in particular has meant that inequality has risen sharply since 1980. South Asia and Europe and Central Asia, are not as homogenous in incomes as North America, but are decidedly more homogenous than East Asia and the Pacific, Sub-Saharan Africa and Latin America. In South

Asia and in Europe and Central Asia, the experience of the last 30 years has been one of increasing inequality within the regions (except for a small decrease after 2005). By contrast, in the initially high inequality regions of Latin America and in Sub-Saharan Africa, inequality has fallen or remained constant. Taken as a whole, we once again see the pattern of income inequality convergence noted previously, this time by regions: in 1980, Gini coefficients ranged from 0.35 (North America) to 0.75 (East Asia Pacific), but by 2013, the range had narrowed to 0.4 to 0.61. In Latin America, inequality rose between 1990 and 2000 before falling subsequently. The period between 1960 and 1990 was, moreover, a more dramatic period of inequality reduction for the region than the 2000s, which have been much vaunted as a time of decreasing inequality in individual Latin American countries.[14]

Figure 3.9: Global Absolute Income Inequality (2005 PPP)
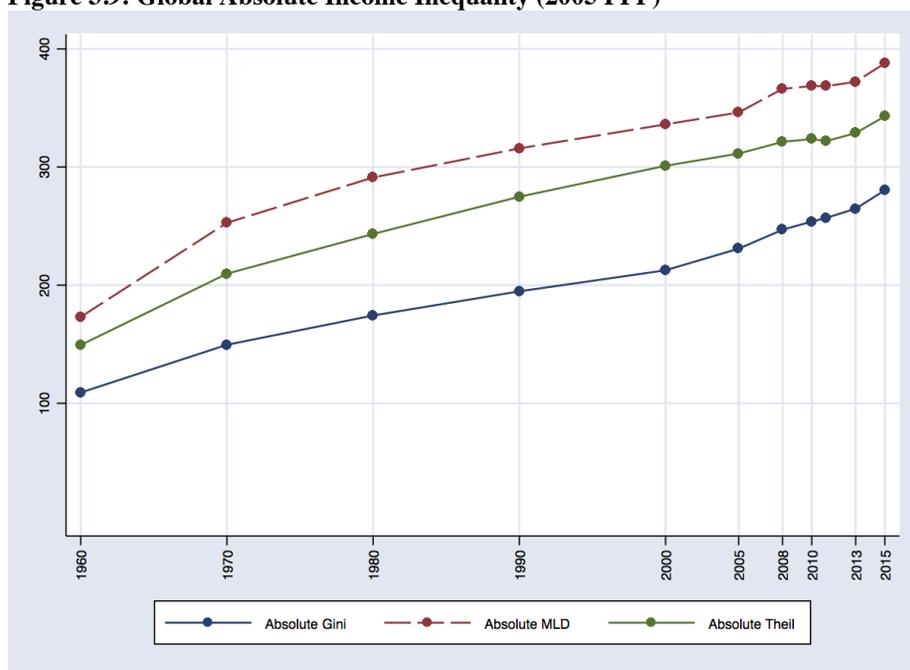
A similar pattern of regional inequality change is also witnessed when we look at other measures of relative inequality such as the mean log deviation (GE(0)), the Theil index (GE(1)), the mean-to-median ratio or the Palma ratio (Table 3.5). Between 1990 and 2013, Latin America and the Caribbean, East Asia and Pacific, and Sub-Saharan Africa experienced sharp declines in all these measures, while South Asia and North America saw increases in these measures.

---

[14] It is important to note the distinction between inter-personal inequality in entire regions (e.g. the Gini coefficient of a region) , as contrasted with the pattern of inter-personal inequality of countries in a region (e.g. the average Gini coefficient of countries in a region). The latter has almost invariably been the focus of studies. The two can, however, evolve differently as differences between as well as within countries matter for the former, unlike the latter.

Table 3.5: Regional Interpersonal Inequality Measures

	Europe and Central Asia	Latin America and Caribbean	Middle East and North Africa	North America	South Asia	Sub-Saharan Africa	World
Gini							
1990	0.47	0.53	0.59	0.38	0.43	0.65	0.72
2000	0.52	0.54	0.6	0.4	0.44	0.64	0.72
2013	0.46	0.5	0.64	0.4	0.44	0.61	0.66
MLD							
1990	0.4	0.53	0.63	0.26	0.31	0.82	1.17
2000	0.58	0.56	0.67	0.28	0.33	0.8	1.13
2013	0.44	0.46	0.8	0.29	0.33	0.71	0.93
Theil							
1990	0.37	0.54	0.68	0.24	0.35	0.89	1.02
2000	0.47	0.56	0.71	0.28	0.36	0.84	1.01
2013	0.36	0.45	0.81	0.28	0.36	0.78	0.82
Mean to Median Ratio							
1990	1.51	1.79	2.27	1.23	1.5	2.64	4.54
2000	1.68	1.83	2.32	1.29	1.5	2.54	4.05
2013	1.39	1.59	2.72	1.29	1.49	2.34	2.9
Absolute Gini							
1990	259.03	137.66	112.17	526.79	21.42	46.55	194.77
2000	309.04	152.93	130.59	633.62	27.29	40.02	212.59
2013	359.49	169.59	225.5	659.09	41.9	52.89	264.68
Absolute MLD							
1990	222.4	137	121.07	361.35	15.46	59	315.75
2000	344.15	157.43	146.78	448.32	20.52	50.34	336.11
2013	342.01	156.07	283.14	473.93	31.66	61.61	372.16
Absolute Theil-T							
1990	205.09	139.6	130.62	331.81	17.48	63.94	274.83
2000	275.26	158.3	154.65	444.27	22.57	52.94	301.05
2013	281.35	153.7	285.32	454.55	34.35	67.28	328.71

Source: Authors' calculations.

Table 3.6 provides comparisons of estimates of global interpersonal relative inequality (Gini coefficient) across a range of different studies. These studies differ in the welfare concept (income, consumption or mixed), data source (e.g. household survey or national accounts means; with or without adjustment of top incomes using administrative/National Accounts data) and exchange rates used (e.g. PPP concepts and base year). Global relative income inequality is estimated to be in a narrow range across the studies and to have been stable in the 1990's with a downward trend after 2000, although the extent of the estimated decline varies.

3.4.　　　Aggregate Analysis: Fast Growing Countries

As we have seen above, there have been marked changes in the world distribution of individual material advantage, in particular in the period after 1990. It may be asked which countries' changes are most responsible for bringing about the observed overall changes in the world distribution? For instance, is there a single country such as China or a group of countries such as the much discussed BRICS (Brazil, Russia, India, China, and South Africa) or another such group of sufficiently large and fast-growing countries that plays this role? To investigate this question, we define a group of 'fast growing countries' and then proceed to examine their impact on the world distribution and to compare this to the impact of the BRICS countries without China, and of China alone. In doing so, we exploit a unique feature of the GCIP, which makes it straightforward to characterize user-defined aggregates of countries.

In recent years much attention has been given to the impact of a few large developing countries that have experienced exceptionally high growth rates on the structure and dynamics of the world economy. During the first decade of the century, the majority of global growth started to be produced in developing countries. This fact has had enormous implications for understanding such matters as which countries drive global consumption and investment dynamics, the location of business opportunities, the changing political status of countries, and how the profile of material affluence, poverty and inequality in the world have been shifting.

What qualifies a country as a 'fast-growing country'? Either an absolute or a relative conception of what it means to be fast-growing could be adopted. However, growth-rates from different sources are quite different. In particular, growth-rates of income from surveys are often very different than those from national income accounts for the very same country-year. We are able to make such a comparison for a larger number of cases as we estimate the means that would have resulted from 'missing' survey data in non-survey years (through interpolation or extrapolation from survey-years) and because we estimate the survey-based estimates that would have resulted from an income survey when we have only a consumption survey. Although the literature has recognized the lack of consistency of national accounts and survey-mean levels and (less frequently) of growth rates (Deaton, 2005; Ravallion, 2003b) it may not have been fully appreciated how pervasive the difficulty is, in particular encompassing developed as well as developing countries, occurring in all major regional groups and extending over all decades for which there is data.

Table 3.6: Comparison of Global Inequality Estimates from Different Studies

	1988	1993	1998	2003	2005	2008	2013
Global Consumption and Income Database (Income)	72.5	72	71.6	70.6	69.4	68	66
Anand and Segal (2014) (with top incomes)	72.6	72.7	72.2	73.5	72.7		
Anand and Segal (2014) (without top incomes, survey means)	70.5	70.7	69.8	71.1	70.1		
Anand and Segal (2014) (without top incomes, Household Consumption from National Account)	73.9	72.1	71.1	70.6	69.8		
Lakner and Milanovic (2015) (National Account Means + top heavy Pareto imputation)	76.3	76.1	77.2	78.1		75.9	
Lakner and Milanovic (2015) (National Account Means)	71.5	70.5	70.6	70.7		67.6	
Lakner and Milanovic (2015) (only survey means)	72.2	71.9	71.5	71.9		70.5	
Milanovic (2012)	67.8	69.3	68.8	70.1			
Milanovic (2005)	61.9	65.2	64.2				
Milanovic (2002)	62.5	65.9					
Bhalla (2002) (Income)	67		65				
Bhalla (2002) (Consumption)	66		63				
Bourguignon & Morrisson (2002)		66					
Chotikapanich et. Al (1997)	65						
Dikhanov & Ward (2002)	69		68				
Dowrick & Akmal (2005) (GK)		64					
Dowrick & Akmal (2005) (Afriat)		71					
Sala-i-Martín (2006)	65	64	64				
Bourguignon (2012)	71		69			66	

Notes: Milanovic (2012): Table 4, p. 14: Gini from row 5 (2005 PPP, sep. rural-urban prices for China, India & Indonesia); Theil from row 3 (2005 PPP, sep. rural-urban prices for China only); 2002 figures for 2003 benchmark. Milanovic (2002): Table 16, p. 72: Using full sample; Table 19, p. 78 (decomposition): Only for common sample. Milanovic (2005): Table 9.4, p. 108: Using full sample; Table 9.5, p. 112 (decomposition): Only for common sample. Bourguignon (2012): Figure 1, only approximate, because read-off from figure; 1988 refers to 1989, 1998 refers to 1997, 2008 refers to 2006. Otherwise: Anand and Segal (2008), Table 1: Survey estimates allocated to benchmark according to rules with micro data: 1988: Bhalla (2002), CVR (1997), and DW (2002) all refer to 1990; 1993: BM (2002) refer to 1992; 1998: Bhalla (2002) refers to 2000, and DW (2002) refers to 1999. GCIP's 1988 estimate refers to 1990.

Source: Klasen, Scholl, Lahoti, Ochmann and Vollmer (2016), Lakner and Milanovic (2013) and additions by authors

To focus on countries that might have a non-negligible effect on world economic dynamics or on the structure of the world-distribution, we isolate ones with a population of ten million or more. As a rough-and-ready criterion, we initially identify among these the top fifteen countries by per-capita income growth rate (see Table 3.7) according to surveys or national accounts. As can be seen, the countries that emerge from these two exercises are rather different. In particular, the survey based classification leads to some surprising results as it leads to the

Table 3.7: Fastest Growing Countries

Country	GDP per Capita Growth Rate	Survey Mean Growth Rate
China	9.01	8.96
Cambodia	5.53	3.43
Vietnam	5.38	6.24
India	4.67	2.93
Sri Lanka*	4.64	2.77
Korea, Rep.*	4.58	2.05
Chile*	3.72	1.46
Thailand	3.71	4.21
Bangladesh*	3.68	2.61
Malaysia*	3.64	2.74
Poland*	3.62	1.74
Mozambique	3.55	2.88
Indonesia	3.54	4.69
Uganda	3.26	4.44
Dominican Republic*	3.07	0.96
	Countries part of top 15 fastest growing countries based on Survey Mean Growth Rate but not part of list based on GDP per capita growth rate	
Kazakhstan	2.50	4.89
Nepal	2.45	4.21
Chad	2.20	4.19
Russian Federation	0.68	3.73
Pakistan	1.77	3.05
Tunisia	3.06	2.91

Note: Only countries with a population of greater than 10 million in 2013 are included in the analysis. The countries marked with * in the table would not be part of the top 15 fastest growing countries if survey mean growth rates were the criteria used instead of GDP per capita growth rate. Compounded annual growth rates are reported in the table.

inclusion of some countries (e.g. Nepal, Chad, Pakistan and Russian Federation among others) not viewed in widespread perceptions (typically based on national accounts based growth rates) as star performers. As survey-based growth rates attributed to individual years are based on interpolation between sometimes distant survey-years they may be more subject to variations resulting from single-year variations or changes in methodology as compared to annual growth rates. However, survey and national accounts based characterizations differ for many reasons (for instance because certain components of national accounts do not enter into household income or consumption at all). We prefer to use both the sources for aspects of the analysis. Specifically, we assess changes to the national economy (through which we identify the countries in the aggregate) in terms of per-capita growth rates from national accounts, and we assess changes to individual material well-being in terms of reported levels of consumption or

income from household surveys. Specifically, we select the fifteen countries with highest per-capita-income growth according to national income accounts and with population of at least ten million as our fast-growing emerging countries, and then proceed to analyze the changes experienced by their population in aggregate as well as the contribution of these changes to those experienced by the world population, over a relevant time period.

Figure 3.10 depicts the initial Gini coefficient and changes in the Gini coefficient between 1990 and 2013 for these fifteen fastest-growing countries. Contrary to the belief that fast growth is generally accompanied by increase in inequality we do not find any such clear pattern here. Six countries experience a decline in inequality, while five experience an increase and four witness no change in inequality during this high growth period.

Figure 3.11 shows the change in distribution for China, the next fourteen fastest growing countries as an aggregate and the rest of the world between 1990 and 2013. The median income for China increased 57 percent during this period, while the rest of the fast-growing countries

Figure 3.10: Change in Gini Coefficient between 1990 and 2013 for the Fastest Growing Countries
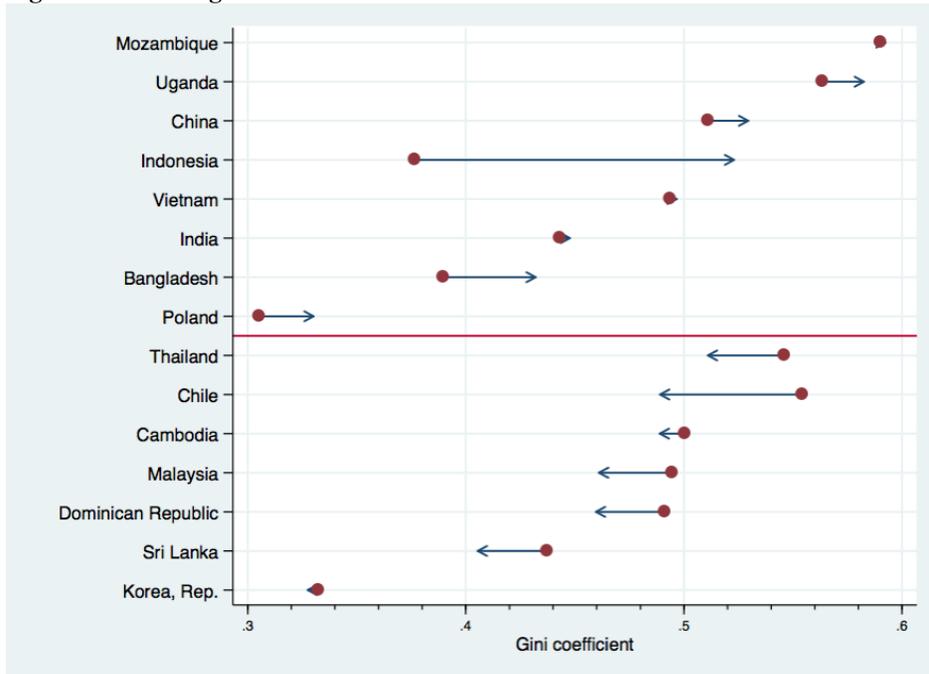
Note: Countries in the list are the top 15 fastest growing countries by GDP per capita during 1990-2013 period with population greater than 10 million.

Figure 3.11: Change in Income Distribution for Fast Growing Countries (2005 PPP)
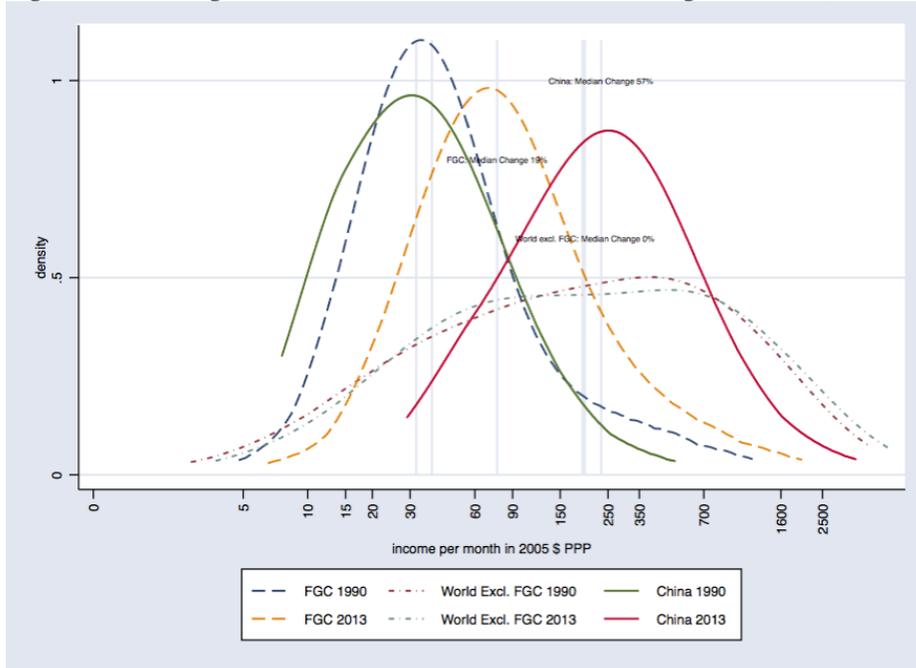
saw an increase of only 19 percent and the rest of the world witnessed no change in median income. The movement of the Chinese distribution and to a lesser extent that of the other fast-growing countries account for almost the entirety of the movement of the world distribution during 1990 to 2013. We undertook a similar exercise for China, BRIS (Brazil, Russia, India and South Africa) and the rest of the world and find that BRIS and rest of the world contribute only marginally to change in global income distribution as compared to China (analysis not shown here).

3.5.   Aggregate 'Social Welfare'

Given the uneven nature of growth as well as the changes in inequality, how might we go about assessing whether the world taken as a whole has experienced economic welfare gains during the period in consideration and if so to what extent? A useful tool in this regard is the Generalized Lorenz Curve (GLC) which allows us effectively to rank distributions in terms of welfare. Shorrocks (1983) showed that for *any* welfare function that is Schur Concave (i.e. responds negatively to regressive Pigou-Dalton Transfers and is therefore inequality averse) and that is positively responsive to income, a given distribution of income would provide more welfare than another distribution if its GLC were everywhere higher. Moreover, all such income-focused

welfare functions would agree on the welfare ranking of two situations if and only if such 'dominance' is established. Figure 3.12 depicts the generalized Lorenz curve for the world distribution of income in four periods of time. Clearly, by any measure, compared to 1960, the welfare of the world had increased by 1990. Perhaps somewhat surprisingly however, there was no unambiguous welfare improvement between the 1990 and 2000 distributions of income, since the generalized Lorenz curves lie almost on top of each other. However, following that period, we see that by 2013, by any measure, global welfare had again increased. Assessment of the actual change in welfare in quantitative terms is also possible but requires the choice of a specific welfare function (or class of functions).

While the GLC provides a framework for welfare comparisons, using a growth incidence curve provides a more detailed depiction of the beneficiaries of growth across this period. Growth has been broadly 'inclusive' in the limited sense that it has taken place across percentiles of the world distribution, but it has been rather uneven across the different percentiles, and the temporal pattern of increases has also varied across percentiles, as shown in Figures 3.13 and 3.14. Between 1960 and 2010, the poorest experienced a greater share of their cumulative growth, in particular of consumption, early in the period. However, we observe an interesting hump-shape (which was first described as an elephant shape by Lakner and Milanovic (2015), arising during the interval of the greatest dynamism (2000 to 2013 and 1990 to 2013)[15]. The middle-income groups – those between the 40th and 60th percentiles - saw their incomes rise rapidly in this period, while those in a rather higher income bracket (80th -95th percentiles) saw their incomes grow much more slowly, mainly because the richer countries in which they disproportionately lived experienced lower growth. Of course, problems of estimation of top-incomes in household surveys provide an essential qualification to any such conclusion. Lakner and Milanovic (2015) plot GICs for the period 1988-2008 and show a larger uptick in growth after the 95th global percentile than observed in the period 1990-2013 in our graphs. For the same period 1988-2008 we also see higher growth above the 95th percentile using GCIP data, but this effect diminishes if the period is extended to 2013. GICs for the longer time frame of 1960-2013 are far flatter than for the more recent period and also consumption GICs show a far more pronounced pattern than the income GICs.

---

[15] The literature that has been spawned by the so-called elephant graph is already sizable. For a critique and response see Corlett (2016) and Lakner and Milanovic (2016)

Figure 3.12: Global Generalized Lorenz Curve (2005 PPP)
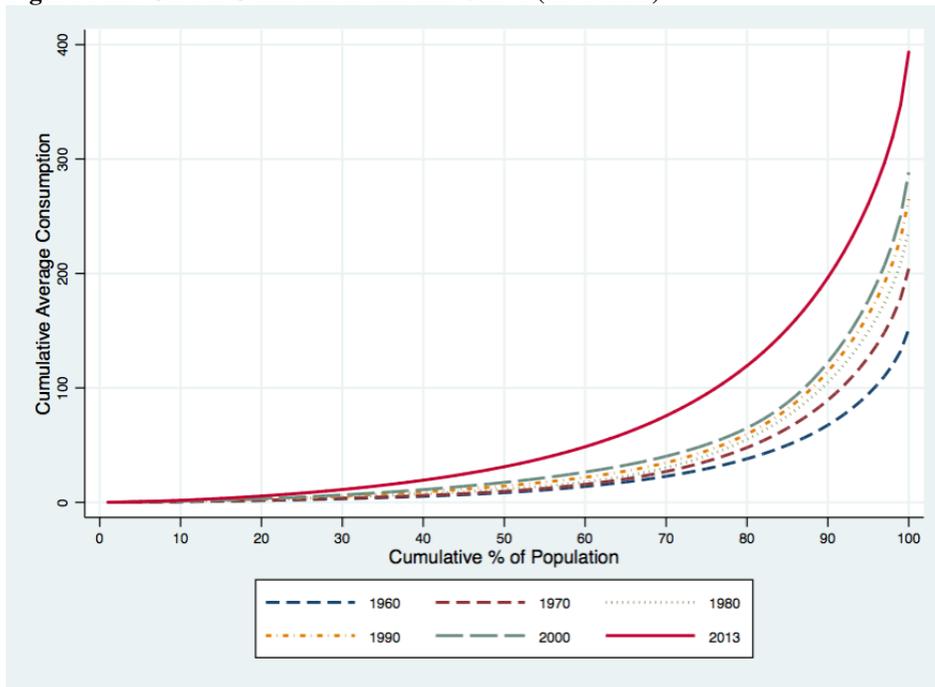
The importance of the Chinese experience in this transformation can be seen to be critical since the Chinese population makes up a disproportionate share of those in the fortieth to seventieth percentile range in 2013. Another way to arrive at this conclusion is to look at the relative

Figure 3.13: Global Consumption Growth Incidence Curve (2005 PPP)
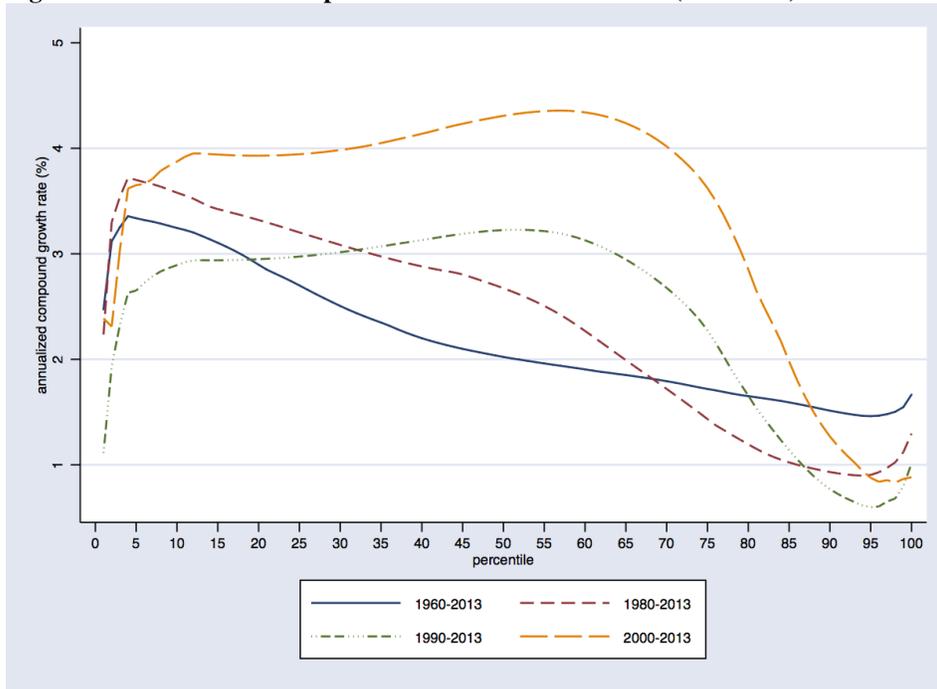
Figure 3.14: Global Income Growth Incidence Curve (2005 PPP)
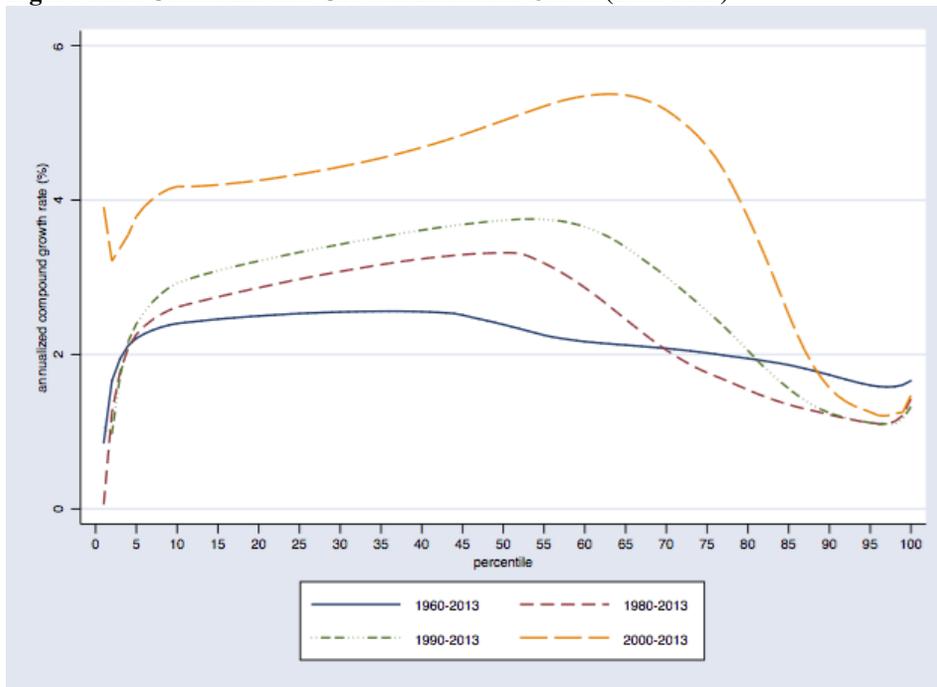
position of the populations of several large countries over time. This is done in Figures 3.15 and 3.16. In 1990, all of the US population enjoyed incomes that would place them in the top quintile of the world income distribution. The Chinese percentile distribution was in the group of bottom countries depicted. By 2013, however, the Chinese percentile distribution dominated those of several of the other countries and had caught up with Brazil. China is now truly the 'Middle Kingdom' of the world as most of its population lie between the fortieth and seventieth percentile of the world distribution. Nigeria is now dominated in relative terms by the other percentile distributions depicted here (although it was not previously). The Chinese relative position looks better when we plot income distributions but is slightly weaker when we plot the consumption distributions of different countries due to the relatively high savings rate of Chinese population. In income distributions,

China and Brazil overlap, but China is substantially lower than Brazil when one plots consumption distribution. Also, even though GDP per capita for Brazil in 2013 was about a third higher than China the survey means are almost the same. This is due to differences in the disparity between national accounts and surveys in the two countries. In

Figure 3.15: Relative Position of Select Countries in 1990 based on Income (2005 PPP)
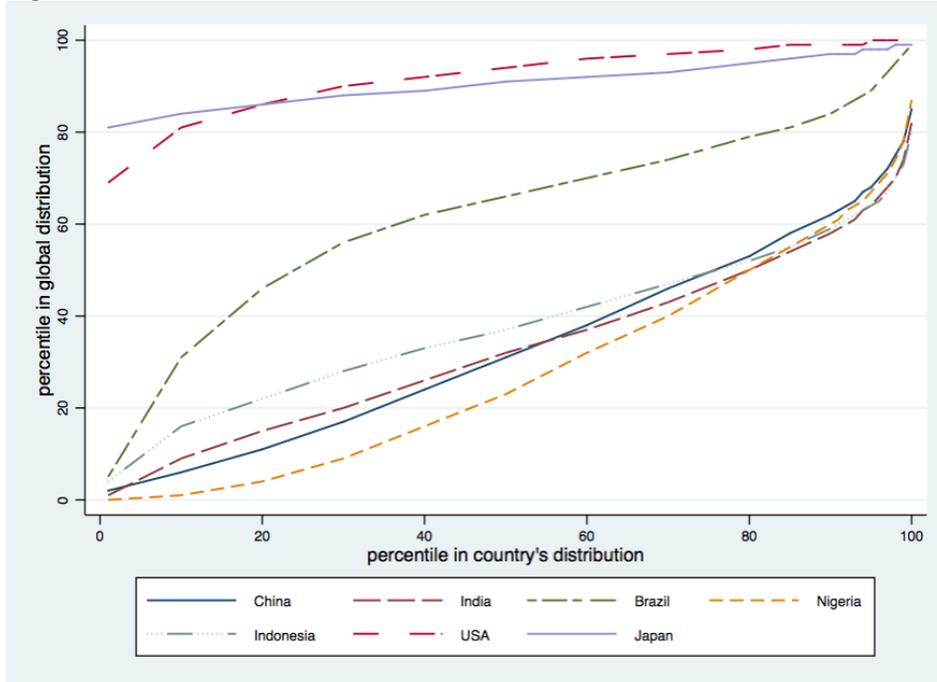
China national accounts and survey means track each other closely whereas they don't to the same extent in Brazil. These differences in bilateral comparisons due to the measures chosen shows that the broader picture too might be dependent on the particular indicator used for comparisons and that systematic exploration of such dependence is necessary to arrive with greater confidence at overall results.

Figure 3.16: Relative Position of Countries in 2013 based on Income (2005 PPP)
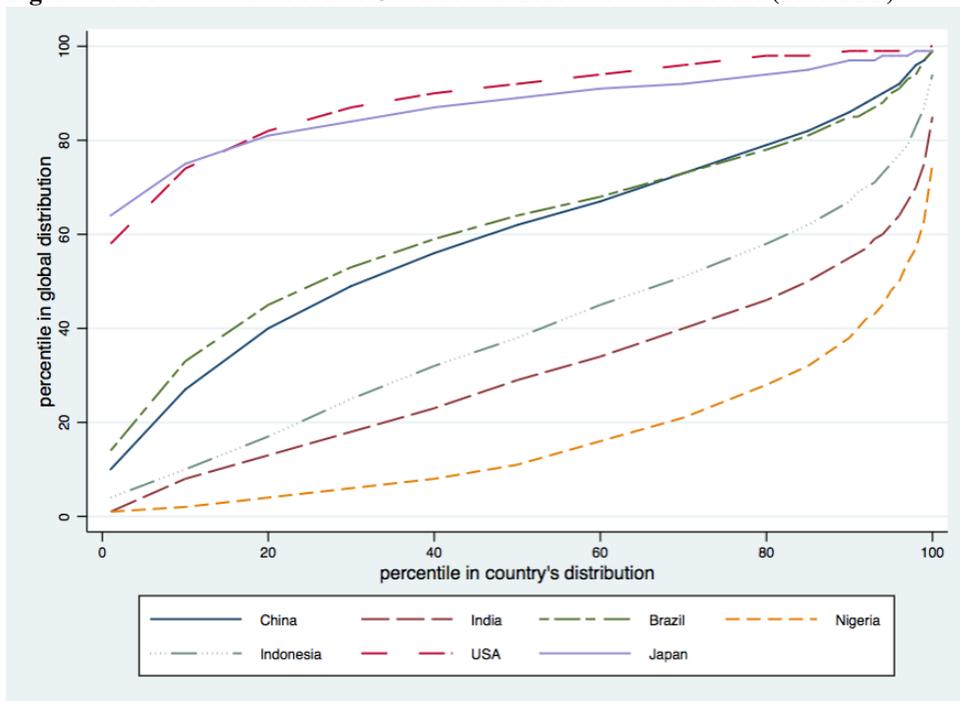
3.6. Sensitivity of the Global Distribution to Alternate Methodological Choices

In constructing the databases and in the subsequent analysis we make several choices. These include: adopting a conversion factor for transforming means expressed in national currency units to a common unit which can then be used for cross-country comparisons, making a decision on whether or not to standardize the distributions by estimating income from consumption (or vice versa), and using means from surveys or national accounts. Our benchmark analysis, on which the discussion above is based, uses 2005 ICP consumption PPPs calculated based on the EKS method to convert national currencies into common international currency units; we

standardize the distributions and use means from surveys (all discussed further in the paper presenting an overview of the database, previously cited). In this section, we modify some of these choices and evaluate the impact. This is in keeping with our larger goal: that the GCIP should be flexible and permit alternate choices so that we might make choices deemed more warranted for specific purposes as well as better understand the robustness of specific conclusions.

We use 2011 ICP consumption PPP's as alternate currency conversion factors in the analysis below. We also construct a consumption database (Mixed Surveys with Income Preference or MSIP) in which we do not attempt to estimate consumption distributions from income surveys but rather pool the distributional information without prior adjustment of either the distributions or the means. There is consumption preference only in the sense that if we have both income and consumption surveys to choose from for a country-year we prefer consumption surveys, which is a preference that could as well have been reversed. We again exclude China from some of the analysis to evaluate the country's impact on the world.

Figure 3.17: Kernel Density Graphs for Various Global Distributions for 2013
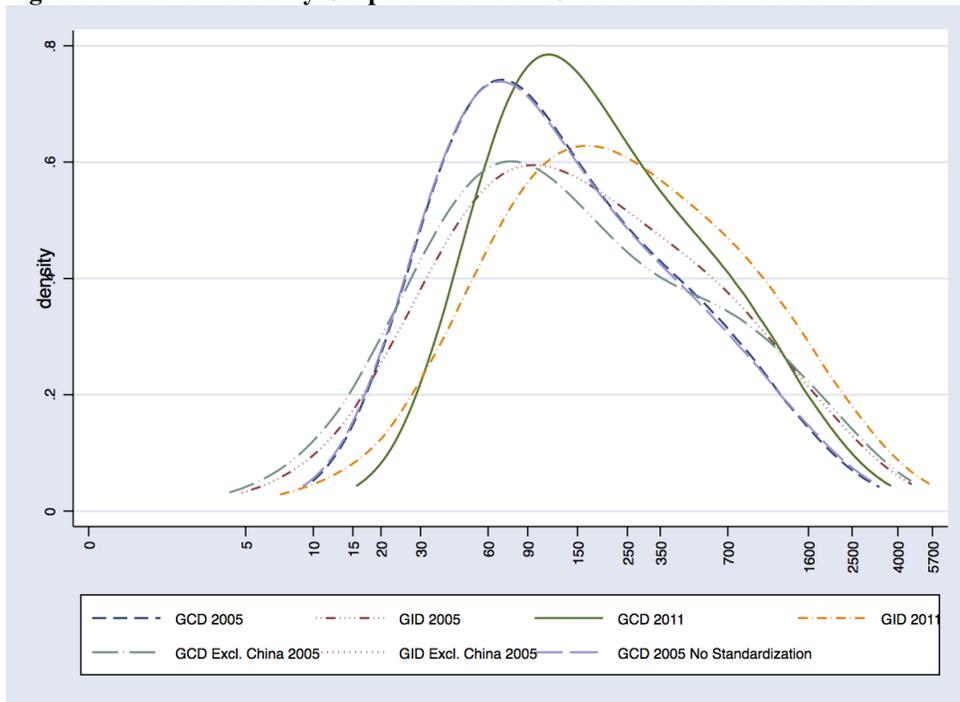
Figure 3.17 depicts the density function corresponding to the resulting alternate databases. The income database (GID) has a seemingly wider effective range than the consumption distribution

Figure 3.18: Global Generalized Lorenz Curve for World Excluding China (2005 PPP)
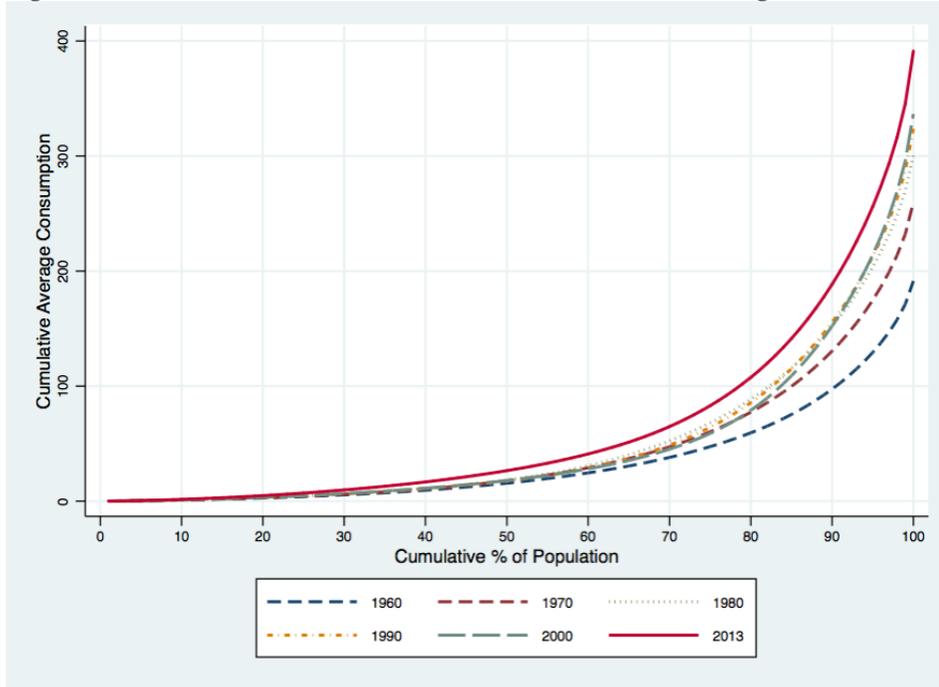
(GCD) for the 2005 PPP base year, apparently because there are many poorer people in the world whose reported or estimated consumption appears to be significantly higher than their income. It is unclear whether this comparison arising from a 'snapshot' of the two distributions reflects a temporary phenomenon or one that is more durable, and why it is present. As we discuss further below, there is evidence that the gap has increased markedly since 1990, suggesting the possibility of growing debt or asset-depletion driven consumption. Interestingly, the density of the income distribution arrived at without any standardization (i.e. by assuming that a consumption distribution is in relative terms exactly the same as an income distribution when only information about the latter is available, but combining this with an estimate of the consumption mean) is almost identical to the consumption distribution. This may be because many populous regions in the world tend to have consumption surveys, in particular in South Asia and Africa and often in East Asia too, and thus provide much of the mass of the density functions in both cases. Although we show the income distribution for the 2011 PPP base year on the same graph, it must be kept in mind that these are in different units (2011 international dollars) that are not strictly comparable to the units in which the 2005 PPP base year distributions are expressed. Nevertheless, interpreted in terms of the perspective of the country with respect to which international dollars are normalized (the US), in which there was marginal domestic inflation in the period in question, it is clear that the distribution shows a shift to the

81

right (i.e. rather higher incomes in the world, especially for the sections of the population in the middle of the distribution) as a result of the application of the 2011 PPPs.

The level of inequality as judged by different inequality measures varies across the alternate datasets (Table 3.8). Trends in global inequality over the past two decades are similar whichever dataset or measure we use. Global inequality appears to increase or remain unchanged between 1990 and 2000 and then to decline between 2000 and 2013. China has a major role to play in the moderate but decided global inequality decline in the period 1990-2013. Excluding China from the world results in ambiguous changes in global income inequality: the 1990-2000 period witnessed small increase or no change in inequality and 2000-2013 period saw a small decline in inequality. Average within-country inequality is 7-10 Gini coefficient points higher when using the income as the welfare concept as opposed to consumption, with the non-standardized version coming close to the average when the consumption database is used (Table 3.9).

The latest round of ICP data was collected in 2011 and a resulting report released in 2014 (World Bank (2014)). The PPP estimates from the 2011 round are quite different from ones based on extrapolations using the 2005 PPPs, perhaps in part due to methodological changes, in part due to sampling variations, and in part due to real changes in the structure of the world economy between the two base years, which are reflected in PPPs that represent data collected in a given year. Most poor countries are estimated to be richer relative to developed countries when 2011 PPP are used than was estimated earlier based on 2005 PPPs. It is important to appreciate that

Table 3.8: Global Relative Interpersonal Inequality for Various Global Distributions

	Gini			MLD			Theil-T		
	1990	2000	2013	1990	2000	2013	1990	2000	2013
GCD 2005 PPP	0.71	0.69	0.64	1.04	0.98	0.79	0.98	0.95	0.77
GID 2005 PPP	0.72	0.72	0.66	1.17	1.13	0.93	1.02	1.01	0.82
GCD 2011 PPP	0.67	0.65	0.59	0.88	0.83	0.66	0.85	0.82	0.64
GID 2011 PPP	0.69	0.68	0.62	1.01	0.98	0.81	0.9	0.9	0.72
GCD No Standardization 2005 PPP	0.71	0.7	0.64	1.06	1	0.8	1	0.97	0.78
GCD Excluding China 2005 PPP	0.68	0.68	0.66	0.97	0.98	0.87	0.86	0.89	0.81
GID Excluding China 2005 PPP	0.7	0.71	0.68	1.12	1.18	1.02	0.91	0.98	0.88
GCD Excluding China 2011 PPP	0.64	0.65	0.62	0.82	0.84	0.74	0.75	0.79	0.71

Source: Authors' calculations.

this applies to comparisons between any pair of country-years in our space-time tableau and not merely to the base years in question. The PPPs for individual consumption by households (used in the benchmark version of GCIP) were revised downwards (on which see e.g. Deaton and Aten, 2017; Ravallion, 2014). These changes would have a sizable impact on global between-country inequality and through that on overall global inter-personal inequality if the 2011 PPPs were used. Our estimates indicate that global between-country inequality in 2013 as estimated by the Gini coefficient was five points lower (0.50 vs 0.55) when using the 2011 PPPs for comparisons rather than the 2005 PPPs (Table 3.10). Global inequality was also lower in 1990 and 2000 when using the 2011 PPPs.

Table 3.9: Within-country Inequality for Various Global Distributions (Gini, population weighted Average)

	1990	2000	2013
GCD	0.36	0.39	0.39
GCD-No Standardization	0.37	0.4	0.4
GID	0.46	0.47	0.46

Source: Authors' calculations.

To test the sensitivity of our welfare results to Chinese growth, we redo the Global Generalized Lorenz curves for the rest of the world excluding China (Figure 3.18). Here we find that, without China, between 1970 and 2000 there were no pareto-improvements in the income distribution (some percentiles in 1970 performed better than in 2000 and also in the intermediate years). The period after 2000 clearly saw large Pareto gains even if one excludes China.

Table 3.10: Between-country Inequality for Various Global Distributions (% of Total Inequality, MLD)

	GCD 2005	GCD 2011	GID 2005	GID 2011	GID Excluding China 2005	GID Excluding China 2011
1960	84	82	76	74	66	61
1970	85	83	78	75	68	64
1980	84	82	76	73	67	63
1990	80	77	67	63	67	62
2000	75	71	63	58	66	61
2005	73	68	59	53	65	59
2010	69	65	56	51	63	58
2013	69	63	55	50	62	57

Source: Authors' calculations.

The cumulative growth record across the percentiles in the global distribution depends on the dataset used to evaluate it (Figure 3.19). For each dataset considered (varying for example according to whether it is based on standardized surveys of a specific type, mixed surveys, or

exchange rate concepts) there was a peak in growth rates between the fiftieth and sixtieth percentile, but the levels of cumulative growth vary tremendously by percentile. It can be seen that the most consequential choice is that between income and consumption based estimates, with income-based estimates showing a bell-shaped growth incidence curve with much higher

Figure 3.19: Global Growth Incidence Curves for Various Global Distributions for 1990-2013
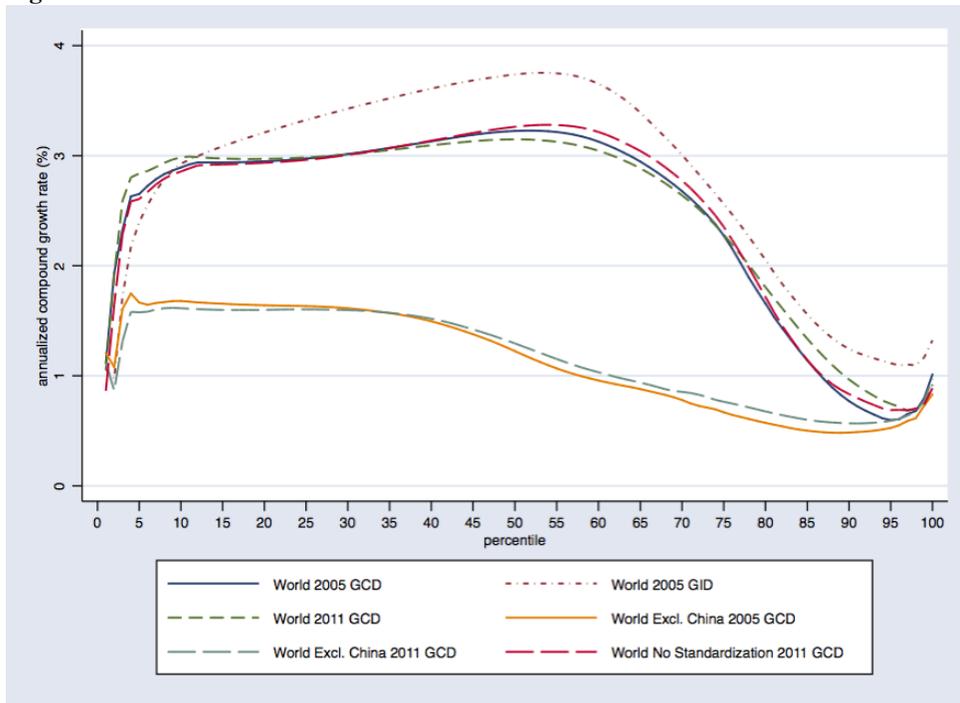
cumulative growth rates for the middle-sections of the world population. The high growth-rates of the middle sections of the population are, however, crucially dependent on the role of China, as can be seen by comparison with the growth incidence curve for income that strips it out. It is also noteworthy that there is slight non-monotonicity as one moves to the right, with the very highest percentiles in the world having somewhat higher cumulative growth rates than those immediately below them. These results could of course be further substantially influenced by the inclusion of top-income information from non-survey sources (an extension to the database we are presently working on). The growth-rates for consumption for the poorer sections of the world population also seem to have been generally considerably greater for consumption than for income over the time-period, suggesting that the magnitude of the excess of consumption over income observed of the relatively poor in recent years may have increased fairly recently.

3.7.   Conclusion

We have in this paper provided a broad overview of the changes in the world distribution of income and consumption over the last fifty-five years. Using a more comprehensive and internally consistent database than previously available, we reproduce some patterns that have emerged in other recent research and drew some new conclusions about the differing patterns in regional distributions and the sources of changes in global inequality as well the overall record of world development. The enormous importance of China for our conclusions in all of these areas of concern stands out. Collective global economic 'welfare' as assessed according to standard assumptions appears to have increased, but the major increase took place in the early 2000s, and may or may not be durable.

We gave some support to previous findings (for example the disappearance of the global 'twin peaks' and the appearance of a degree of 'inequality convergence') and identified new stylized facts (for example the fact that the Sub-Saharan African countries' inequality, when measured in a way so as to standardize these surveys with those from other countries, is among the very highest in the world, and in particular higher than in Latin America). We identify some new findings of interest for further research (such as that the lowest percentiles of the world distribution appear to be dissaving). Moreover, we show that much of our picture of what has happened to the world depends on choices with respect to the variable depicting individual advantage, the choice with respect to exchange rate (e.g. PPP of a given base year). Each of these may be justified depending upon the purpose at hand, but has significant implications for our understanding of how the world's population, as a whole and in its parts, has been faring.

GCIP is an invaluable resource to study evolution of material living conditions in the world. In Reddy and Lahoti (2016), we used GCIP data to show that while absolute poverty has declined by most measures, but by comparison to a higher and plausible poverty line, poverty across the world has not changed significantly over three or even five decades. In Jayadev, Lahoti and Reddy (2015), we discuss the various ways of defining the global middle class and use GCIP to study the evolution of the middle class.

4. How Serious is the Neglect of Intra-Household Inequality in Multi-dimensional Poverty Indices?[1]

Abstract

Income-based as well as most existing multidimensional poverty indices (MPI) assume equal distribution within the household and thus are likely to lead to yield a biased assessment of individual poverty, and poverty by age or gender. In this paper we first show that the direction of the bias depends on how these measures use individual data to determine the poverty status of households, while the impact of these assumptions on inequality between individuals cannot be determined a priori. We then use data from the 2012 Indian Human Development Survey to create a standard household-based MPI closely related to that proposed by Alkire and Santos (2014) as well as UNDP (2014), and compare it to an individual level MPI that individualizes education and nutrition and some aspects of the living standards dimensions. We find that the poverty rate of females is 14 percentage points higher than that of men in our individual MPI measure but only 2 percentage points higher when using the household-based measure. Similarly, the age differentials in poverty are much larger using the individual-based measure. Using a decomposable inequality measure, we find the contribution of intra-household inequality to the total inequality in the individual deprivation score inequality to be 30% and total inequality is also some 30% *higher* using the individual-based measure, while inequality among the poor is found to be 5% *smaller* using the individual measure.

4.1. Introduction

The ultimate objective of measuring poverty and inequality is to determine the wellbeing of individuals. But most empirical analyses of poverty assumes that resources are distributed equally, or according to need, within the household and equate poverty status of the household with the poverty status of all individuals in the household. Such household-based poverty measures are then often used to track trends in well-being, target social and economic programs, and measure the impact of interventions.

But the assumptions of equal or needs-based distribution is inconsistent with the theoretical literature on intra-household bargaining, which has shown that well-being outcomes depend on

the bargaining power within the household where equal distribution would be more of the exception than the rule. These bargaining models have received overwhelming empirical support in the literature (e.g. Lundberg & Pollak, 1993; Lundberg, Pollak, & Wales, 1997; Manser & Brown, 1980; McElroy & Horney, 1981; Grossbard-Shechtman 1993; Gersbach & Haller, 2001; Edlund & Korn 2002, Chiappori 1988, 1992).

More generally, there is overwhelming evidence collected across multiple contexts over the last two decades on intrahousehold inequalities against the need-based or equal distribution assumption (e.g. Haddad et al., 1997; Quisumbing & Maluccio 2000; Aronsson, Daunfeldt, & Wikström, 2001; Alderman, Chiappori, Haddad, Hoddinott, & Kanbur, 1995). In particular, substantial and consequential gender inequalities in the allocation of resources have been shown to exist in many contexts, with particular sizable gaps existing in some regions of the developing world, particularly parts of South Asia and the Middle East (e.g. World Bank, 2011; Klasen & Wink, 2002; 2003, Asfaw, Klasen, & Lamanna, 2010; Rosenzweig & Schultz, 1982; Hazarika, 2000).

As a result of this it is likely to be the case that household-based assessments of poverty by gender understate the gender gap in poverty, at least in some parts of the developing world.[2] And similarly, often-done analyses of child poverty or poverty among the elderly will yield biased results as the equal distribution assumption is unlikely to hold (e.g. Dreze & Srinivasan, 1997; Corak, Fertig & Tamm, 2008; Deaton & Paxson, 1998). More generally, poverty rates might be biased and their distribution by region or household type distorted, leading to biased assessments of individual well-being and policies, and biased targeting.

Even though this has been long recognized there have been only few attempts at measuring poverty and inequality using truly individual level achievements. The dominant approaches in both unidimensional income and multi-dimensional poverty measurement (MPI) use the household as the unit of analysis to determine poverty status of individuals.

In 1990, Haddad and Kanbur assessed how serious the neglect of intra-household distribution is when considering poverty in calorie intake (Haddad & Kanbur, 1990). Using Philippine data they show that 30 to 40 percent of all inequality is intra-household inequality and would be missed if individual data was ignored. They also find that ranking between males and females

---

[2] At the same time, there have also been some unverified claims about gender gaps in poverty, such as the widely made claim in the 1990s that 70% of the world's income poor are female. If one assumes equal distribution at the household level, it is impossible to arrive at such a figure; but since no information existed on the actual unequal distribution of poverty, this number was a pure conjecture. See Marcoux (1998) for a discussion.

reverses when using individual data, with poverty rates among women being higher when using some poverty measures.

Several methods have been developed in recent years that allow one to estimate intra-household inequality using only household-level monetary information. (Lise & Seitz, 2011, Chiappori, Fortin & Lacroix, 2002, Dunbar, Lewbel & Pendakur, 2013, Browning, Chiappori & Lewbel 2013). Case & Deaton (2002) and Chiappori & Meghir (2014) provide an excellent review of the various approaches used in the literature. But first and foremost they note the serious challenges when doing so. This is due to the presence of public goods within the household, the difficulty in identifying the sharing rule within the household given limited data and varying preferences across household members; and lack of sufficient data on individual consumption and time use for household members all of which complicates the estimation of intra-household inequality. Case & Deaton (2002) in their review conclude that most methods in the literature rely on controversial, easily challenged and non-transparent assumptions. Also none of these methods have gained widespread acceptance. They suggest that the best way forward might be collection of more data on individual consumption and interviewing multiple people in the household. But even such improved data will not solve the conceptual problem of determining how the use of household-specific public goods can be attributed to its members.

In contrast to the income dimension where household-specific public goods make an assessment of individual income poverty particularly difficult (e.g. Klasen, 2007), many non-income deprivations, e.g. in the health and education dimensions, can, in principle, be attributed to individuals so that an individual multidimensional poverty measure appears more feasible at first sight. And in fact, these individual-level data are typically available in standard survey instruments. Yet in existing popular multi-dimensional poverty measures such as the Multidimensional Poverty Index (MPI) used by UNDP and OPHI (see Duclos 2011; Alkire & Foster, 2011; Alkire & Santos, 2014, d'Ambrosio & Chakravorty, 2010), deprivations are also determined at the household level; and all individuals within the household are assigned the deprivation and poverty status of the household without any differentiation within the household. In cases where the household is deemed deprived or poor then all individuals in the household are deemed to suffer equally from these deprivations. Consequently, the gender or age-segregated poverty numbers obtained from these calculations are unreliable at best, and deeply misleading at worst. And even overall poverty numbers, trends, and correlates might be similarly affected.

An individual might be personally deprived in several or most dimensions to be actually deemed poor; but the household might be deemed non poor based on combined assessment of household members and hence the individual though deprived would be deemed as non-poor or vice versa in a household-based analysis. The bias this generates in household-based multidimensional poverty assessments depends on how the thresholds for household poverty in a dimension are set, or how the individual-level data is combined to create a household-level indicator. The deprivation thresholds can be defined in a *restrictive* way where the achievement of the worst-off member of the household has to be above the threshold for the household to be non-deprived. In these cases the deprivation rates among individuals would be *higher* in the household-based analysis, as long as not all households are indeed equally deprived in that dimension, than the analysis which takes into account only individual's personal status. But deprivation thresholds could also be defined in an *expansive* way, where only the achievement of the best-off individual has to be above the threshold for the household to be non-deprived. In such cases, the deprivation rates among individuals would be *lower* in the household-based analysis if not all are as well off as the best-off as compared to an individual-based analysis. UNDP and OPHI's MPI use a mix of indicator threshold definitions – restrictive and expansive – so that the net bias of their neglect of intra-household inequality is not clear a priori.

There have been some survey-based multidimensional measures proposed exclusively for different demographic groups within the population (Alkire et al., 2013; Roche, 2013), but most focus only on a subset of the population like women or children. Ramaya, Lahoti & Swaminathan (2014) construct an individual level multidimensional poverty measure for adults in Karnataka, India. They found substantial gender differences in poverty that are absent when using household measures. Bessell (2015) proposes an individual deprivation measure for adults based on custom-made surveys in the Philippines, finding rather little gender inequality in this deprivation measure. While these studies are instructive, they are only focused on particular groups and thus cannot assess poverty at the individual level for the entire population or assess to what extent household-based analyses under- or overstate individual poverty and inequality. Also, they are based on particularly detailed, unique, and often custom-made surveys using small samples, making replication at higher scales and across contexts difficult (and costly).

To our knowledge, ours is the first paper to present a multi-dimensional poverty measure at the individual level that accounts for intra-household inequality across the entire population. We aim to ascertain if individual welfare and inequality among them is under- or over-stated when using household-based analysis as compared to an analysis based on individual level data. Using data

from India, we use our measure to estimate individual poverty and inequality as well as the size of the bias of household-based analyses.

In our application, we find that women and older individuals in India are far more deprived and poor than men and younger individuals. This simple fact is obscured and gender and generational differences are absent when measuring poverty and inequality using the standard household-based approach. In particular, the poverty rate of females is higher by 14 percentage points than men in our individual MPI measure but only 2 percentage points higher when using the household-based measure. The poverty rate among individuals aged fifty and over is higher by 46 percentage points than among children aged between 7 and 18 years of age in the individual measure, compared to only 2 percentage points when using the household-based measure. Using a decomposable inequality measure, we find the contribution of intrahousehold inequality to the total inequality in the individual deprivation score inequality to be 30% and total inequality is also some 30% higher using the individual-based measure, while inequality among the poor is found to be 5% smaller using the individual measure. We also find that in over 60 percent of households the average deprivation level of women in the household is greater than average deprivation level of men.

At the same time, we note that our approach to individualize poverty measurement can only be seen as a first attempt in this direction and is hampered by insufficient data on individual well-being in standard household surveys; improved data would likely lead to even larger differentials in poverty by age and gender, at least in a country such as India. The paper is organized as follows. In the next section we discuss our theoretical framework, section three discusses data and methods, section four results, and section five concludes.

4.2.   Theoretical Framework

We adapt the theoretical framework for assessing the impact of neglecting intra-household differences in the uni-dimensional setting presented in Haddad & Kanbur (1990) to a multi-dimensional setting.

Let's assume that wellbeing of individuals is measured by y. In a uni-dimensional setting, wellbeing is generally measured by consumption, income or nutrition. In a multi-dimensional scenario let d (≥ 2) represent the number of dimensions in which well-being is assessed and y_{ij} (≥ 0) represent the achievement of individual i in dimension j. Let the total number of individuals be N (i = 1,2,….N) belonging to H households (h=1,2…H). The dimensions used in

multi-dimensional poverty analysis commonly include education, health, and indicators of standard of living. Each dimension j is assigned a weight w_j. The weights represent the relative importance assigned to each dimension by the analyst. Let z_j denote the threshold below which an individual is deemed deprived in dimension j, and let z be the row vector of dimension thresholds. For each individual i, let g_i^0 denote the deprivation vector of d elements, whose elements g_{ij}^0 are defined by $g_{ij}^0 = w_j$ when $y_{ij} < z_j$, while $g_{ij}^0 = 0$ otherwise. We assume that information on individual's deprivation with respect to any particular dimension is binary i.e. 1 if deprived and 0 if non-deprived.

As discussed above, in household-based multidimensional poverty assessments thresholds are not defined based on achievements of each individual but collectively for the household, denoted by z_j^h. All members of the household then are assumed to have the identical deprivation vector g^{0h}. We can construct a weighted deprivation count vector c, whose entry for the ith individual is the sum of the weights for the dimensions in which the individual is deprived, $c_i = \sum_{i=1}^d g_{ij}^0$. When using household data, the deprivation score for all individuals in the household is identical and given by $c^h = \sum_{i=1}^d g_j^{0h}$. The difference between c and c^h for individuals within and across households and their distribution is the main object of interest in this paper.

Are the levels of c and c^h systematically different, and are individuals of certain groupings favored to have higher well-being in one over the other? Do the differences in c and c^h impact multi-dimensional poverty analysis? The answers to these questions depend on how the underlying dimension thresholds are defined in the household-based analysis and the extent of within-household disparity in achievements for the dimensions and the poverty line.

Household deprivation thresholds can be defined in various ways. For some indicators, there exists only a household-based indicator and the implicit assumption is that, in this dimension, we are dealing with is at least partly household-specific public good, even though in some cases these might be excludable and/or rivalrous. This is, for example, the case the standard of living dimensions of UNDP's MPI that examine electricity and water and sanitation access for the household, or the ownership of durable goods to determine household-level deprivation in these dimensions. In these cases, individual data is not available on these household-specific goods. While of course one cannot be sure that all household members profit equally from access to these public goods (esp. use of some of the durable goods might be quite unequal), it is very hard and information-intensive to assess intra-household inequality in access to these public goods.

As a result, most surveys do not contain individual-level information on these dimensions. We will return to this issue in the empirical assessment below.

More important for our purposes here, however, is that for some dimensions, household-level assessments and thresholds are built up from individual-level data that *is* available in the surveys. We classify the most commonly used thresholds of using individual-level data to assess household-level deprivation into two types, restrictive and expansive.

The deprivation threshold is said to be *restrictive* when the achievement of the *least* well-off person (or the overwhelming majority of members) has to be above the threshold for the household to be non-deprived. For example, in UNDP's MPI, the threshold that deems the entire household to be deprived in nutrition if any one member of the household is undernourished, is such a *restrictive* one. This could generally be represented by a deprivation function defined as $g_j^{0h} = w_j \ if \ \min(y_j^h) \leq z_j^h \ and \ 0 \ otherwise$.

In such instances the average value of the deprivation score across the population for the dimension would be higher than if individual data were used to define deprivations i.e. $\mu(c_j^h) \geq \mu(c_j)$, as long as there is no perfect equality among all household members in this dimension. In other words, the number of individuals with the deprivation would be *higher* using household thresholds than individual-specific thresholds. Individuals within the household who are better off would be deemed deprived due to deprivation of the worst-off household members. In the Indian example below, men on average are better off than women in most well-being-dimensions, and would more likely be misidentified as deprived in such dimensions.

The deprivation threshold is said to be *expansive* when the achievement of only one (or a minority of individuals) has to be above the threshold for all individuals in the household to be non-deprived. For example, in the MPI, the entire household is deemed non-deprived in education if at least one household member has five years of education. This can be generally represented as $g_j^{0h} = w_j \ if \ \max(y_j^h) \leq z_j^h \ and \ 0 \ otherwise$. In such instances the average value of the deprivation score across the population for the dimension would be higher than if individual data were used instead to define deprivations i.e. $\mu(c_j^h) \leq \mu(c_j)$. The number of individuals with the deprivation would be *lower* using a household-based assessment than using

individual-specific thresholds. For example, women in India, who on average are worse off than men, are likely to be misidentified as non-deprived using such thresholds.[3]

In theory, a household deprivation threshold can be constructed from individual data using a linear or some other combination of data on individual household members. For example, one could deem the entire household as deprived in education if the average number of years of education for all adult members of the household is below five years. In such instances individuals below the average achievement of the household might have lower deprivation score whereas individuals above the average of the household might have higher deprivation scores than when individual data is directly used. But in the standard household MPI measures none of the thresholds are defined in this way.

The extent of the disparity in individual deprivation status within the household in each dimension would determine amount of under- or overstatement. For example, if the within household disparity in nutrition deprivation, which is defined in a restrictive way, is large and many households have only one undernourished person while others in the household are not undernourished then the extent of overstatement would be large. On the other hand, if within household disparity is small and in deprived households most members are undernourished then the extent of overstatement would be small. An analogous argument can be made in case of more expansive deprivation thresholds. The aggregate impact of various dimensions on difference between c and c^h would depend also on the type of thresholds, weighting and to what extend each deprivation misidentifies individuals. In most popular multi-dimensional measures some indicators are defined restrictively while others are defined expansively, so some of the over- and understatement of deprivation rates would lead to opposing biases and thus partially cancel each other in the aggregate measure. For example, in the MPI proposed by OPHI and UNDP, the educational achievement dimension is defined in an expansive way, while the educational enrolment, and the undernutrition dimensions are defined in a restrictive way.

To create an aggregate measure of the incidence of multi-dimensional poverty based on these dimensional deprivation data, a recently proposed influential approach by Alkire and Foster (2011) is to select a cutoff value `k' and any individual with a weighted deprivation score above `k' is considered multidimensional poor i.e. $\rho_k(y_i, k) = 1 \ if \ c_i \geq k \ and \ \rho_k(y_i, k) = 0 \ if \ c_i < k$ where ρ_k is the identification function. For aggregating poverty over the population one simple

---

[3] The household-based MPI uses third methods to assess household-level deprivation based on individual-level achievement. In the mortality indicator, a household is deemed deprived if a child has died in the last 10 years. We use the same procedure for that component in the individual MPI.

approach is to measure the percentage of population that is poor. The headcount can be formally defined as $H(y) = q/n$ where $q = \sum_{i=1}^{n} \rho_k(y_i, k)$ is the number of persons who are identified as poor. UNDP's MPI has, for example, adopted this approach for identification and aggregation.

The impact of differences between c and c^h on the poverty headcount or any of the other poverty measures depends on the distribution of deprivation scores, especially with respect to k.

Does the pattern of poverty incidence across groups change when taking intra-household differences into account? Suppose we divide the population into m mutually exclusive groups population proportion δ_l with l =1…m. The multidimensional headcount of poverty can be written as

$$H(y) = \sum_{l=1}^{m} \delta_l H_l(y); H(y^h) = \sum_{l=1}^{m} \delta_l H_l(y^h)$$

and the contribution of a group to poverty in the two cases can be shown to be

$$C_l(y) = \frac{\delta_l H_l(y)}{H(y)}; C_l(y^h) = \frac{\delta_l H_l(y^h)}{H(y^h)}$$

For the contributions to poverty to be different the intra-household inequality has to be very different across the two groups, similar to the uni-dimensional case (Haddad and Kanbur 1990). If we divide the population by regions and find that intrahousehold inequality is substantially higher in rural areas as compared to urban regions and most of the dimension thresholds are defined in expansive way (leading to understatement of poverty in household measure), then the rural contribution to poverty might increase when using individual data. In such cases there is a possibility of change in poverty rankings across the two regions i.e. if $H_U(y^h) > H_R(y^h)$ then with substantially higher intra-household inequality we can get $H_R(y) > H_U(y)$.

What about inequality in deprivation scores? Is the distribution of the total deprivation score of individual deprivation scores c more or less unequal than compared to when deprivation scores are based on a household-level assessment c^h? And how do the intra-household and inter-household components of inequality change? In the uni-dimensional case Haddad and Kanbur (1990) show that inequality (using all Lorenz-consistent inequality measures) is understated when using household-level data. The individual level c can be seen as the result of a mean-preserving spread that is bound to increase inequality.

This is, however, not always true in the multidimensional case. In the MPI measures that do household-based analysis all individuals within the household are assigned the same status and hence intra-household inequality is assumed to be zero by definition and all inequality is inter-household. So the intra-household inequality is underestimated. But when moving from a household-based assessment to an individual assessment, *inter*-household inequality is also affected. The change in inter-household inequality depends on the distribution of deprived individuals across the households in the population.

If the deprivation thresholds are restrictive and deprived individuals are concentrated in some households so that households either have all deprived individuals or no deprived individuals, then the inter-household inequality is the same when using household and individual data. In simple terms, if there is no misidentification of deprivation status of individuals when using household thresholds, then total inequality is the same in household and individual analysis. But if deprived individuals are spread more widely across households so that deprived and non-deprived individuals live in the same household then inter-household inequality could be under- or over-stated by a household-level assessment (that deems everyone deprived as long as a single individual is deprived) depending on the exact dispersion of deprived individuals across households, with examples for under- and overestimation provided in the footnote.[4] Similar considerations hold when the expansive definition is used so that again it is an empirical question whether total inequality is higher in an individually-based or a household-based assessment of multidimensional deprivation.

What about the pattern of inequality across groups? How does between and within inequality among groups change with use of household vs. individual data? For illustration and in our empirical analysis we use an inequality measure based on a positive multiple of the variance proposed by Seth & Alkire (2014). This is an absolute inequality measure as opposed to a relative inequality measure more commonly used in assessing income inequality. An absolute inequality measure is invariant to any additive changes to deprivation scores, while a relative inequality measure is invariant to any proportional changes to deprivation scores. As discussed in

---

[4] Consider a small hypothetical population consisting of two households (A and B) each consisting of three members. We assess MPI poverty for this population based on achievements in five dimensions, which have restrictive thresholds and are equally weighted. In the first scenario (S1) each member of household A is deprived in one indicator and one member of household B is deprived in four indicators while other two members are deprived in one and three indicators, respectively. In the household assessment all members of household A will have deprivation score of 0.2 and due to restrictive thresholds, all members of household B will have deprivation score of 0.8, resulting in a Gini coefficient of 0.3. In the individual assessment the six individuals will have deprivation scores of 0.2, 0.2, 0.2, 0.2, 0.6 and 0.8 and the Gini of the individual deprivation score will be higher at 0.32. Let consider another plausible scenario (S2) with two members of household A with no deprivations and the third with deprivation in one dimension, and each member of household B is deprived in four dimensions. In the household assessment of MPI the deprivation scores and Gini of deprivation score remains the same as scenario S1, due to restrictive thresholds. But in the individual assessment the Gini of deprivation scores decreases to 0.17. So the exact dispersion of deprivation scores across individuals and households determines whether inequality is under- or over- estimated.

Seth & Alkire (2014), a relative inequality measure for counting based poverty measurement approaches will provide contradictory conclusions depending on whether one measures attainment or deprivations, which is not the case in absolute measures. Also each deprivation has a direct or intrinsic importance justifying the normative assessment of inequality in absolute distances (see also Klasen, 2008). The measure can be expressed as

$$V(c) = \frac{4}{n} \sum_{i=1}^{n} [c_i - \mu(c)]^2$$

This can decomposed into within-group and between-group components as below:

$$V(c) = \frac{4}{n} \sum_{l=1}^{m} \sum_{i=1}^{t^l} [c_i(k)^l - A^l]^2 + 4 \sum_{l=1}^{m} \delta_l (A^l - A)^2$$

The first part of the equation is the total within-group component of total inequality. It can be viewed as a positive multiple of the sum of variance within each group. A^l is the average deprivation score in the group l i.e. $A^l = \mu(c^l)$ and A is the average deprivation across the entire population. The second component is the total between group inequalities. The between group contribution is defined as

$$C_B(c) = 4 \frac{\sum_{l=1}^{m} \delta_l (A^l - A)^2}{V(c)}$$

The within group component is just $1 - C_B(c)$. If we are using household level thresholds and data, then

$$V(c^h) = \frac{4}{t} \sum_{l=1}^{m} \sum_{i=1}^{t^l} [c_i^h(k)^l - A^{hl}]^2 + 4 \sum_{l=1}^{m} \delta_l (A^{hl} - A^h)^2$$

$$C_B^h(c) = 4 \frac{\sum_{l=1}^{m} \delta_l (A^{hl} - A^h)^2}{V(c^h)}$$

As discussed earlier the difference in $V(c)$ and $V(c^h)$ depends on the type of deprivation thresholds and distribution of deprived individuals across households. So the direction or extent

97

of difference in the between-inequality contribution by the two methods, and its impact on group-based inequalities, remains an empirical question.

In sum, household-based assessments of multidimensional will provide a biased account of individual multidimensional poverty. It will then also bias the assessment of poverty by groups as well as the measured total inequality in deprivations. While for some definitions of household-based assessments and levels of intra-household inequality in deprivation, one can assess the sign of the bias, for others this is not possible a priori and will essentially become an empirical exercise to which we now turn.

4.3. Data and Methodology

We use data from the 2012 Indian Human Development Survey (IHDS) to construct the multidimensional poverty measures. IHDS is a nationally representative, multi-topic panel survey of 42,152 households across India covering all Indian States. Most of the households interviewed were part of an earlier round of IHDS survey in 2005. IHDS covers a wide range of topics, which include health, education, employment, economic status, marriage, fertility, gender relations and social capital. The survey also asked a few sex-disaggregated time-use questions about common household chores like collecting water and cooking. Unlike most household surveys, IHDS recorded individual level asset ownership information for land and principal residence. This individual level information allows one, for example, to calculate gender asset gap within households (see Lahoti, Suchitra & Swaminathan, 2016)

4.3.1. Dimensions and Indicators

We construct a Household MPI measure (henceforth called Global Household MPI), which is based on the Global Multidimensional Poverty Index developed by the Oxford Poverty and Human Development Initiative (OPHI) and also used, in a slightly amended version, by UNDP (2014). An individual MPI measure is constructed using the same dimensions as the Global Household MPI, but by directly measuring individual achievements in some dimensions as opposed to household-level deprivation indicators. The individual MPI measure uses slightly different indicators than the household measure. In order to facilitate ease of comparisons we construct a second household MPI measure which uses exactly the same indicators as the individual MPI. This new household MPI measure is referred to just as the comparable household MPI. All the MPI measures incorporate education, health and standard of living as the

three dimensions. A list of the various indicators, deprivation thresholds and weights used in each of the measures is presented in Table 4.1.

Education

The education dimension is commonly seen as a central capability Sen's and Nussbaum's versions of the capability approach (e.g. Nussbaum, 2003; Sen, 1998). It also impacts other capabilities such as future employment opportunities, self-confidence and the ability to participate in public and political life.[5] There is near-universal acceptance of the importance of education in measuring multidimensional poverty and it is used in all the major MPI measures.

Table 4.1: Dimensions, Indicators and Weights for various MPI measures

Dimension	Indicator	Deprived if ..	Weight
Global Household MPI			
Education	**Years of Schooling**	No member of Household completed 5 years of education	1/6
	Children's school enrollment	One or more HH members between 7 to 15 years of age not enrolled in school	1/6
Health	**Nutrition**	One or more adult HH member is underweight, or any child is undernourished	1/6
	Mortality Among Children	One or more children born to interviewed women in the household died after birth	1/6
Standard of Living	**Electricity**	No access to electricity at home	1/18
	Floor	House Floor made of mud	1/18
	Sanitation	No access to private toilet	1/18
	Water	No access to Safe Water Source within 15 minutes one-way distance from the residence	1/18
	Cooking Stove	HH uses open fire or traditional chulha without chimney to cook	1/18
	Consumer Durables	Owns less than 2 of list of assets (TV, Phone, cycle, refrigerator, motorized vehicle)	1/18
Comparable Household MPI			
Education	**Years of Schooling**	No member of Household has completed 5 years of education	1/3
Health	**Same as Global Household MPI**		
Standard of Living	**Electricity**	No access to electricity at home	1/18
	Floor	House Floor made of mud	1/18
	Sanitation	No access to private toilet	1/18
	Access to Safe Water	No access to Safe Water Source within 15 minutes one-way distance from the residence	1/36
	Water Collection Time	Time taken to collect water by all household members is one hour or more	1/36
	Cooking Stove	HH uses open fire or traditional chulha without chimney to cook	1/18
	Consumer Durables	Owns less than 2 of list of assets (TV, Phone, cycle, refrigerator, motorized vehicle)	1/18
Individual MPI			

[5] A recent example is a law passed in the Indian states of Rajasthan and Haryana prohibiting anyone without certain minimum years of education to contest local level elections.

Education	Years of Schooling	Not completed threshold years of education**	1/3
	Nutrition	Individual is malnourished*	1/6
Health	Mortality Among Children	One or more children born to interviewed women in the household died after birth	1/6
	Electricity	No access to electricity at home	1/18
	Floor	House floor made of mud	1/18
	Sanitation	No access to private toilet	1/18
	Access to Safe Water	No access to Safe Water Source within 15 minutes one-way distance from the residence	1/36
Standard of Living	Water Collection Time	Time taken to collect water by individual's group in the Household is one hour or more	1/36
	Cooking Stove	HH uses open fire or traditional chulha without chimney to cook	1/36
	Cooking Time	person does most of the cooking with unsafe stove	1/36
	Consumer Durables	Household owns less than 2 of list of assets (TV, Phone, cycle, refrigerator, motorized vehicle)	1/18

* Adult ages for 18 years or older is undernourished if BMI is less than 18.5. Individuals 6 to 17 years of age are malnourished if BMI-for-age is two or more standard deviations below the median of the reference population. Children between 0 to 5 years of age are deemed malnourished if their weight for height is two or more standard deviations below the median of the reference population.)

** The threshold years of education for individuals 12 or more years of age is 5 years of education, for children between 7 to 11 years of age is the age-adjusted years of schooling so that they complete five years of education by age of 12. Children below 7 years of age are deprived if half or more household members 12 or more years of age have not completed 5 years of education.

The indicators used to measure education in the global household MPI are proximate literacy and children's enrollment in school. Basu and Foster (1998) argued that presence of a literate person provides positive externality for the entire household. In the global MPI a household with at least one member who has completed 5 years of education is considered non-deprived (UNDP's MPI sets the threshold at 6 years but assumes the same externality). This is an expansive threshold and would lead to underestimation of deprivation rate for this indicator. Despite this externality, education provides first and foremost a benefit to the person who possesses it so that an individual perspective seems warranted. In addition, Ramaya, Lahoti & Swaminathan (2014) argue that differences in literacy among household members might impact power dynamics. So in the individual MPI measure we measure education separately for each person in the household. We deem an individual above 12 years of age as deprived if she/he has not completed five years of education. For children below age 12, we use a different procedure that we outline presently.

The Global MPI uses children's enrollment[6] as a second indicator for education. The right to education is a central right of the international Convention on the Rights of the Child. Since the passage of the Right to Education (RTE) Act in 2009, education is also recognized as a

---

[6] In addition to enrollment IHDS also measures children's achievement in reading, arithmetic and writing. In 2005-06 more than 42% of children enrolled in grade 5 could not read a simple story, indicating the poor quality of schooling (Desai et al., 2010). We do not include this in our MPI measures to maintain comparability with the Global MPI, but ideally this should be part of measuring poverty as enrollment (even though high) is not indicative of actual schooling.

fundamental right of every child in India. A child not enrolled in school indicates acute distress and curtailment of opportunities for that child.

In the individual MPI measure we do not use this indicator (as there would be no equivalent indicator for adults and children outside of this age window). Instead, children between the ages of 7 to 12 are deemed deprived if they have not completed the expected age-adjusted years of schooling. The expected age-adjusted years of schooling is calculated so that children should be on track by age of 12 years[7] to complete five years of education. Since children below 7 years of age have not started schooling, we have no information on them for the schooling indicator. In these cases, we use information on schooling status of other household members as proxy for their potential status. Specifically, children below seven years of age are deemed deprived in education in the individual measure if half or more of household members 12 or more years of age have not completed five years of education. We also test, in later sections, robustness of our results to modifying the schooling threshold for children below seven.

Health

Health is another central capability. The capability refers to being able to have good health, including reproductive health; to be adequately nourished; and to have adequate shelter (Nussbaum, 2003). It also directly and indirectly impacts other capabilities – malnourishment reduces the ability of children to learn (UNICEF 1998), limits ability to participate in social life (e.g. disability) and might increase the need for material resources (Rippin, 2012).

We use nutrition and child mortality as the two indicators for health dimension, the same as the ones used by the Global MPI[8]. Nutrition is an especially important indicator for India given the overall poor state of nutrition (both among children and adults) in the country (e.g. Klasen, 2008). In the household measures, an individual is considered deprived in the nutrition indicator if any of the adult household members (18 years or more) for whom data is collected are underweight (have a BMI less than 18.5), or if any children are malnourished. Individuals 6 to 17 years of age are undernourished if their BMI-for-age is two or more standard deviations below

---

[7] The age of joining school in India is 6 years. So children are expected to complete five year of education by 11 years of age. We provide a buffer of one year to account for later entry into schools. (Dotter and Klasen, 2014)

[8] We had also experimented with using major morbidity and disability as health indicators. One issue with using major morbidity and disability as indicators for health is that they are partially subjective. As Sen (2002) has argued, self-reported health status maybe seriously limited by the respondent's social experience. A respondent brought up in a community with woeful medical and educational facilities might take certain symptoms as "normal" when they clinically might be symptoms of major morbidity. A second issue is that the elderly experience several of the diseases defined as major morbidity, even though we have limited the list to exclude life style diseases like heart disease. Life expectancy among the poor is less than among richer respondents and hence rich are more likely to live long to experience these diseases. We find that individuals in higher consumption quantiles are slightly more likely to be deprived in these two indicators than individuals in poorer consumption quantiles. Hence, we dropped these two indicators from our MPI measure.

the median of the reference population. Children between 0 to 5 years of age are deemed deprived if their weight for height is two or more standard deviations below the median of the reference population[9]. We use the reference population defined by WHO to calculate the nutritional status for children.

For the individual MPI measure, we use individual data whenever available to define nutrition deprivation. If weight or height data is not collected for an individual then the status of the group to which the person belongs defines her or his status.[10] Each age group 0-5, 6-17 and 18 and above are divided into two based on sex. If half or more individuals in the household from the group are nutritionally deprived (based on individual data) then the group is deemed as deprived. All individuals in the group for whom nutritional data was not collected get the status of the group.[11] In cases where data is not collected for any individuals of a particular group then all the members of the group within the household get the status of the age group. The age group status is deprived if half or more individuals in the age group (male or female) are deprived based on individual data. In the previous example if nutritional data is not collected on any of the 5 adult male individuals then they are all deprived if half or more of the adult females for whom data is collected are deprived. Lastly, if data is not collected for an age-group then all individuals for that age group within the household are deprived if half or more individuals in the household for whom data was collected are deprived.

All individuals in the household are considered deprived in child mortality, if the interviewed women in the household report one or more child deaths. There is no difference between the household MPI and individual MPI deprivation status for this indicator. Unfortunately, we do not have other reliable individual level health indicators to replace the child mortality indicator.

Standard of Living

---

[9] IHDS data on birth history providing the exact age in months has not been released yet, only age in years is available. Using age in years would yield inaccurate measure of weight-for-age but this should not have large impact weight-for-height measure. The levels of the two would certainly be different (wasting vs. underweight) and the extent of gender bias might differ, but the bias is likely to be in the same direction. Given the limitation of data, for now, this seems like best measure for 0-5 age group.

[10] These imputation procedures have little impact on deprivation scores in these dimensions, including gender gaps. Specifically, BMI information is collected directly for 60 percent of all adults. Among those with direct observations on BMI 22 percent of males females are deprived in nutrition. In the overall adult sample, after using our assumptions to impute information on adults for whom BMI was not collected, still 22 percent of males and females are deprived in nutrition. The differences in gender bias in nutrition deprivation for children under six for whom we have direct measurement of weight and height (for 69 percent of all children below age of 6) and the entire population below six years of age (after imputing nutrition status based on their group's status for children) is very small (1% vs. 0.10%). Our assumptions for imputing status of children between 6 and seventeen years of age with no information on nutrition reduces the gender bias (that is in favor of girls) as compared to those with direct nutrition information. We have individual nutrition information on 78 percent of all children in 6-17 age group, and 23 percent of boys and 16 percent of girls are deemed deprived in nutrition. For all children 6 to 17 years of age the corresponding numbers are 24 percent and 20 percent.

[11] For example, if a household has five adult males, and individual nutritional data is collected for three adult men, then the HH adult male group is deprived if two or more of the individuals for whom data is collected are deprived. The two individuals for whom data was not collected get the status of the group and the other three have status based on their own individual data.

The standard of living dimension partly captures 'control over one's environment' central capability. It represents some of the material means necessary to achieve many of the basic capabilities. We use the same indicators for our Global household MPI as used by Alkire & Foster (2011). As discussed above, several of the goods are public in nature within the household. This makes it difficult to determine individual ownership or differential access to these goods. Hence for the individual measure we assume living standards as public goods accessible equally by everyone within the household, similar to the assumption made by Ramya, Lahoti & Swaminathan (2014). At the same time, IHDS collects sex disaggregated time use data on some of these indicators and we include those separately in our individual MPI measure. For example, if adult women or men as a group spend more than an hour collecting water daily then that group is deemed additionally deprived. To maintain the same overall weight on the standard of living dimension with the addition of indicators, we lower the weight on the household indicator for the same living standard to accommodate the time-use individual indicator. The comparable household MPI also includes the household time use indicator for water collection. A household is deprived in the indicator if the collective time spent by all household members is greater than one hour.[12]

Other possible dimensions that could be included are empowerment, physical safety and subjective wellbeing among others. In many of these dimensions there is documented gender disparity in favor of men, particularly in countries such as India (Ramya, Lahoti & Swaminathan, 2014; Bessell, 2015) We do not include these so as to maintain comparability with the OPHI Global MPI and also most surveys lack data for measuring these dimensions. Excluding these will result in understating the gender disparity in poverty.

4.3.2. Weighting

We follow the Global MPI in adopting an equal weighting approach across dimensions. Various studies have looked into the impact of alternative weighting schemes on multidimensional poverty measures (Alkire & Santos, 2014, Decancq, Van Ootegem & Verhofstadt, 2013). But since our main goal is to investigate the gender and generational disparity in poverty we adopt the most common and simplest weighting scheme. Each of the three dimensions is assigned an equal weight of one-third and all indicators within the dimensions are weighted equally, except for cases where time-use indicators are used for a standard of living measure. Each standard of

---

[12] We also initially included asset ownership indicators (land and residence) in the MPI measures as done by Ramya, Lahoti & Swaminathan (2014). But these indicators do not necessarily reflect deprivation but are more correlated with location. Urban households are less likely to own land or residence as compared to rural households, but that is not necessarily a result of deprivation in urban areas.

living indicator gets a weight of 1/18, but if a time-use indicator is used then the each of the indicators for that aspect of standard of living gets a weight of 1/36. For example, the safe water indicator gets a weight of 1/18 in the Global Household MPI, but since time use for water collection is also used as an indicator in individual and comparative household MPI the weight is halved – the safe water indicator gets a weight of 1/36 and the time-use indicator gets a weight of 1/36. In robustness analysis we vary the weighting structure across the three dimensions to test our results to different weights.

Households without eligible population

Dotter and Klasen (2014) discuss the various approaches to deal with so-called ineligible populations in the MPI. Several of the indicators used in MPI measure refer to the achievement status of a particular group within the household. For example, the child mortality indicator refers to death of a child in a household of a woman of reproductive age. But as pointed in Dotter and Klasen (2014), if a household never had children then the household cannot suffer from this deprivation. In cases with no eligible population for a particular indicator in a household, the Global MPI assumes that the household is non-deprived in that indicator. This reduces the possibility of a household with missing eligible population from being deemed multidimensionally poor. A household which has never had children would be non-deprived by definition in children's enrollment indicator and child mortality indicator, making it less likely to be judged as poor even if it is deprived in some other indicators. Based on possible solutions to this suggested by Dotter and Klasen (2014) we substitute the missing indicator with an indicator from the same dimension, i.e. substitute the nutrition indicator for the child mortality indicator for a household that never had any children. This would double the weight on the nutrition indicator for those households.[13]

4.3.3. Limitations

With our approach we succeed in at least partly individualizing our multidimensional poverty measure. But we also note that, due to difficult conceptual issues as well as data limitations, we face some challenges. The most serious conceptual challenge is the adequate treatment of children in the health and education dimensions where we need to rely on comparisons with other groups; we thoroughly investigate the impact of these choices on our individual

---

[13] This is not without problems. It assumes that nutrition indicator is substitute for child mortality indicator and also we can no longer decompose the MPI measure by indicators. But given the lack of data on any other equivalent indicator for the missing information this is a reasonable compromise. We can still decompose MPI by dimensions. In cases where there is no information on any of the indicators within a dimension we chose to drop the household from the sample. We find only a small number of instances of this in our data; hence it does not impact the representativeness of the sample.

multidimensional MPI. The most serious data limitation relates to the health dimension where we need to rely on a household-level mortality indicator and do not have a reliably individual health indicator beyond nutrition. Finally, by mostly relying on the household-level information in the standard of living dimension, we may underestimate inequality in access and use of household-specific public goods, including particularly also durable goods and assets.

4.4.   Poverty and Inequality Measures

For both household MPI's, each individual within the household is assigned the deprivation status of the household. The multi-dimensional poverty index M_0 developed by Alkire & Foster (2011) is the sum of the weighted deprivations suffered by the poor divided by the maximum possible number of deprivations. It can be expressed as

$$M_0 = \frac{\sum_{i=1}^{n} \sum_{j=1}^{d} g_{ij}^0(k)}{nd}$$

MPI though simple, has a problem in identification of the poor. Both the usual identification methods – union and intersection – are not usable in this approach. According to the union approach, an individual is deemed poor if she/he is deprived in at least one indicator. This leads to unreasonably high poverty headcounts – on order of more than 85%. As discussed in Dotter & Klasen (2014) the headcount ratios are not only difficult 'political' sell but also the approach is very sensitive to measurement error or cases where indicators do not cover the deprivations well. On the other hand, the intersection method deems an individual poor only if she/he is deprived in all indicators. This leads to unreasonably low poverty rates.

Alkire & Foster (2011) suggested a compromise between the two methods – the dual-cutoff approach. A person is considered poor only if the sum of weighted deprivations for the person exceeds an additional cutoff k. In case of OPHI Global MPI the value of k is 0.33. So any person whose weighted deprivation count is greater than one third is considered poor. But the choice of k is completely arbitrary and the poverty levels and comparisons across regions change for different values of k. Also the use of dual-cutoff leads to assumptions about correlation between poverty indicators that are difficult to justify. Poverty indicators are perfect substitutes below the cutoff k but perfect complements if the weighted sum of deprivations is more than k (Rippin, 2012).

The MPI measure can be decomposed into the product of (censored) headcount (H) or poverty incidence and (censored) average poverty intensity (A):

$$M_0 = \frac{q}{n} \frac{\sum_{i=1}^{n} c_i(k)}{qd} = H.A$$

where q is the number of poor i.e. number of individuals for whom the sum of weighted deprivations is above the cutoff k.

As Rippin (2012) notes this decomposition has two issues: first the censored poverty intensity is dependent on the arbitrary cutoff k; and second, the MPI measure is not sensitive to changes in inequality among the poor. Amartya Sen (1976) defined as one of the properties of good poverty measures to be decomposable into three components: poverty incidence, intensity and inequality. The MPI measure does not satisfy this property. A regressive transfer between two poor individuals that makes the poorer of the two individuals more deprived and the other less deprived but still poor will not cause the MPI measure to change. An inequality sensitive poverty measure should increase upon such a regressive transfer. There are a few counting-based measures which define poverty measures so that they are sensitive to inequality (Bossert, Chakravarty & D Ambrosio, 2009, Jayaraj & Subramanian, 2009, Rippin 2012). A critique of these measures has been that they appear to lose the property of factor decomposability and are difficult to interpret (Seth & Alkire, 2014).

Rippin (2012) proposes a measure that in the identification step assigns different degrees of poverty to a household while using the union method to identify the poor. The degree of poverty is calculated by taking the square of the weighted share of deprivations suffered by the household.

$$CSPI = \frac{\sum_{i=1}^{n} [\sum_{j=1}^{d} g_{ij}^0]^2}{nd}$$

Similar to FGT2 measure in uni-dimensional case, this takes into account inequality in the distribution of deprivations. Any regressive transfer among the poor would now result in an increase in the poverty measure since the additional deprivation to a highly deprived household counts as more than to a less deprived household. The Correlation-Sensitive Poverty Index can be decomposed into the three I's (Jenkins & Lambert, 1997): incidence, intensity and inequality while still being factor decomposable.

$$CSPI = \frac{q}{n}\left[\frac{\sum_{i=1}^{n} c_i}{qd}\right]^2 + 2\left(\frac{\frac{1}{2q}\sum_{i=1}^{n} c_i}{\frac{1}{qd}\sum_{i=1}^{n} c_i}\right) = HA^2(1 + 2GE)$$

We use a variety of measures in our empirical application. To identify the poor we use two different methods –union and dual-cutoff approach. In the union approach an individual is deemed poor if she or he is deprived in any one indicator. In the dual-cutoff approach we use a cutoff of 33 percent as per the methodology of the Global MPI. We test the robustness of our results to changing values of cutoff. The union and dual-cutoff leads to vastly different poverty headcount ratios using the two approaches. An individual whose weighted deprivation count is 0.25 would be deemed poor using the union approach but not considered poor using the dual-cutoff approach. For our headline measure we use the incidence obtained using the dual-cutoff approach. We calculate MPI using the dual cutoff approach and CSPI measure that uses the union approach. MPI and CSPI are decomposed into their various components to investigate incidence, intensity and inequality among the poor. We also use the variance based inequality measure to study within and between group inequalities.

4.5.   Results

We first analyze the deprivation levels by sex and age group (Table 1) in the various indicators before delving into the poverty and inequality measures. As the schooling indicator used in household MPI is defined in an expansive way, we should find that the individual deprivation level should be higher than the deprivation level when using household data. Indeed, we find that 26 percent more individuals are deprived when using the individual data. As hypothesized, adult women, who are the worst-off group, are more likely to be misclassified as non-deprived when household data is used. Access to schooling for women was very poor in India up until very recently, which is reflected in the higher gender differential and also overall higher level of deprivation in the adult age group, and particularly high deprivation among the oldest age group (50+). There is no gender differential in schooling achievement among children in the age group 7 to 18 age group. This is likely because of concerted push in the last decade by the government to increase school enrollment and the passage of right to education act, which makes education compulsory for this age group. The higher level of deprivation among the below 7 age group children is because their deprivation status is determined by older members of the household, who tend to be more deprived; but note that there is no gender gap here suggesting that our

method does not impute a gender gap into the data. We conduct robustness analysis to test the sensitiveness of our results to changing deprivation level for this age group. As expected, the gender and the age differentials are substantially reduced when using the household measure.

The household nutrition indicator is defined in a restrictive way, with all household members considered deprived if any adult or child is undernourished. As predicted, this results in the household measure indicating substantially higher overall deprivation levels (25 percentage points) than the individual nutrition indicator. There is no significant gender differential in the nutrition indicator among adults. Among the age group 6 to 17 years boys are slightly more deprived than girls. Across all age groups, between 20 to 25 percent of the individuals are undernourished. The level for children below six years of age is lower because we are using the weight-for-height indicator (wasting) instead of the more common weight-for-age measure (underweight)[14]. According to NFHS-3 conducted in 2005-06 about 20 percent of children below five were undernourished as per the weight-for-height measure, but just over 40 percent were undernourished if we use the weight-for-age measure (Table 10.1, Page 270, IIPS 2007)[15]. But there is no substantial gender differential reported for both the measures and hence, even though our levels are lower, the conclusion on gender differentials won't be impacted due to the use of weight-for-height measure[16].

Several household standard of living indicators have witnessed declines in deprivation rates as compared to data reported in National Family Health Survey in 2006 (see Table 1, Alkire and Seth (2015) for more details). Electricity deprivation has reduced from 33 percent to 17 percent, safe water deprivation from 16 to 10 percent, and sanitation deprivation from 70 to 63 percent among the directly comparable indicators. The time use indicators indicate the extra burden on women of not having access to basic amenities. 39 percent of adult women below 50 years of age are directly impacted due to smoke from unclean cooking stoves, while none of the men of this age group suffer directly as a result of unsafe cooking stoves in the household (Ezzati & Kammen, 2002). Women are also more likely to spend time collecting water from outside the household. Better data on time use and access to household public standard of living resources among household members is needed to differentiate their impact on individuals in the household. The household-level time use indicator on water collection gives an incorrect picture of no gender differential in time spent which can be discerned from the individual time use

---

[14] This is due to lack of data on the exact age (in days) of children from IHDS-2. These data will likely be released later and then this can be updated.
[15] These estimates broadly match the estimates from IHDS-1 (Gaiha, R et al., 2010)
[16] 20.5% of boys and 19.1% of girls below five years of age have weight-for-height below 2 standard deviations of the median of the reference population and the corresponding numbers for weight-for-age are 41.9% and 43.1% according to NFHS-3.

indicator. Even though we do not account for it, lack of access to private toilets is also likely to impact women more severely than men. Diane Coffey et al. (2015) find in a survey of Indian villages in some northern and central states that women are more likely to use toilets if they have access to one than men, indicating the greater need among women for better sanitation facilities.

Table 4.2: Proportion of individuals deprived in various indicators

Deprivation	Male 0-6	Female 0-6	Male 7-18	Female 7-18	Male 19-50	Female 19-50	Male 50+	Female 50+	Total
Education									
Household Schooling	0.15	0.16	0.08	0.08	0.07	0.08	0.14	0.17	0.1
Individual Schooling	0.48	0.5	0.14	0.14	0.23	0.42	0.49	0.77	0.36
Children Enrollment	0.06	0.07	0.08	0.09	0.07	0.06	0.07	0.06	0.07
Health									
HH Nutrition	0.53	0.53	0.52	0.53	0.44	0.45	0.41	0.42	0.47
Individual Nutrition	0.2	0.2	0.25	0.22	0.21	0.22	0.23	0.23	0.22
Child Mortality	0.21	0.23	0.23	0.25	0.19	0.2	0.21	0.19	0.21
Standard of Living									
Access to Electricity	0.22	0.22	0.19	0.2	0.14	0.14	0.15	0.15	0.17
Type of House Floor	0.48	0.49	0.46	0.46	0.36	0.37	0.36	0.36	0.4
Access to Sanitation	0.7	0.71	0.68	0.68	0.61	0.61	0.59	0.6	0.63
Access to Safe Water	0.1	0.1	0.1	0.1	0.09	0.1	0.1	0.1	0.1
HH Water Collection Time	0.16	0.15	0.18	0.19	0.14	0.13	0.12	0.11	0.15
Individual Water Collection Time	0.01	0.01	0.02	0.04	0.03	0.07	0.03	0.06	0.04
Access to Clean Cooking Stove	0.66	0.67	0.64	0.64	0.56	0.56	0.56	0.56	0.59
Time with unclean Cooking Stove	0	0	0	0.05	0	0.39	0.01	0.21	0.12
Consumer Durables	0.27	0.29	0.22	0.24	0.18	0.19	0.23	0.27	0.22

Table 4.3 presents the multi-dimensional poverty measures for the three different MPI definitions. The Global Household MPI measure is directly comparable to the Global MPI constructed by OPHI. Comparing our estimates to OPHI estimates based on 2005-06 NFHS survey points to substantial reductions in all poverty multi-dimensional measures. The headcount ratio has declined from 54 percent to 40 percent and the MPI from 0.283 to 0.19. The

comparable household MPI[17] measure uses only the schooling indicator for education and adds a household water collection time use indicator to make it more comparable with the individual MPI measure. There is no change in the MPI between Global and Comparable measures, only minor changes in the headcount and the poverty intensity. Headcount and MPI are slightly higher for females across age group in these two household-based measures. But using individual data the MPI is substantially higher among adult women than men; as intimated in Table 4.2, this is mostly related to higher education deprivation among adult women. The higher MPI is mostly driven by differences in headcount ratios as poverty intensity varies only slightly across gender and age groupings (see also Dotter and Klasen, 2014). There is no gender differential among children in the MPI which is a promising development and driven by lack of education differentials between boys and girls. The overall levels of poverty are higher when using individual as compared to household data. This suggests that the expansive definition used for education is more important for the overall MPI than the restrictive definition used for nutrition. This is not really surprising given the strong age-dependence of educational deprivation. Even households where many people are uneducated will often have one young person with at least 5 years of education, making the entire household non-deprived in the household-based MPI. This clearly shows the problems associated with such an expansive definition.

Table 4.3: Comparison of various multi-dimensional poverty indicators constructed using dual-cutoff approach (k=33%) by age-sex categories

	Male 0-6	Female 0-6	Male 7-18	Female 7-18	Male 19-50	Female 19-50	Male 50+	Female 50+	Total
Global Household MPI									
Headcount	0.47	0.49	0.44	0.46	0.34	0.36	0.39	0.41	0.4
Intensity	0.48	0.5	0.47	0.47	0.47	0.47	0.52	0.52	0.48
MPI	0.23	0.24	0.21	0.22	0.16	0.17	0.2	0.21	0.19
Comparable Household MPI									
Headcount	0.46	0.48	0.42	0.43	0.33	0.35	0.38	0.4	0.38
Intensity	0.51	0.52	0.48	0.49	0.48	0.48	0.52	0.53	0.49
MPI	0.24	0.25	0.2	0.21	0.16	0.17	0.2	0.21	0.19
Individual MPI									
Headcount	0.54	0.56	0.29	0.29	0.33	0.48	0.54	0.78	0.44
Intensity	0.55	0.56	0.5	0.5	0.52	0.55	0.55	0.54	0.53
MPI	0.3	0.31	0.14	0.15	0.17	0.27	0.3	0.42	0.24

Note: Headcount is measured as the proportion of population that is deemed poor.

The MPI is sensitive to the arbitrary poverty cutoff `k', and is not sensitive to inequality among the poor. To address this, we present results for CSPI (Table 4.4) that takes inequality into account and uses the union method for identification of the poor. The headcount for the union

---

[17] The comparable household MPI measure is designed so that it uses the same indicators as the individual measure and is directly comparable.

method is above 85 percent, as anyone deprived in one indicator is considered poor. There is very little variation in headcount across gender-age groupings but the poverty intensity is higher among adult women. Overall inequality in deprivation scores among the poor is higher in the individual than in household MPI, but inequality among poor adult women is less than among poor adult men. Overall the CSPI among adult women is substantially higher than adult men, reflecting the pattern found in MPI analysis.

We calculate absolute inequality measures in deprivation scores across the entire population and among the poor using the measures proposed by Seth and Alkire (2014) and described in section 2 (Table 4.5). Total inequality is decomposed into within and between components for various socio-economic groups. For India we find that inequality in deprivation scores is higher by about 30% when using individual rather than household data. Also of the total inequality 30 percent of inequality in individual MPI is due to intra-household disparity. This component of inequality, which is the focus of this paper, is totally absent from household MPI by definition.

Table 4.4: Comparison of various multi-dimensional poverty indicators and measures constructed using union approach by age-sex categories

	Male 0-6	Female 0-6	Male 7-18	Female 7-18	Male 19-50	Female 19-50	Male 50+	Female 50+	Total
Global Household MPI									
Headcount	0.91	0.91	0.9	0.9	0.85	0.85	0.83	0.84	0.87
Intensity	0.31	0.32	0.29	0.29	0.24	0.25	0.27	0.28	0.27
Inequality	0.3	0.3	0.31	0.31	0.48	0.45	0.56	0.53	0.42
CSPI	0.14	0.15	0.12	0.13	0.1	0.1	0.13	0.14	0.12
Comparable Household MPI									
Headcount	0.91	0.91	0.9	0.9	0.85	0.85	0.83	0.84	0.87
Intensity	0.31	0.33	0.29	0.29	0.24	0.25	0.27	0.28	0.27
Inequality	0.33	0.33	0.34	0.34	0.5	0.48	0.58	0.54	0.44
CSPI	0.15	0.16	0.12	0.13	0.1	0.11	0.13	0.14	0.12
Individual MPI									
Headcount	0.9	0.9	0.88	0.88	0.82	0.84	0.83	0.9	0.86
Intensity	0.35	0.36	0.24	0.24	0.24	0.32	0.33	0.44	0.3
Inequality	0.35	0.34	0.47	0.48	0.66	0.5	0.49	0.24	0.49
CSPI	0.18	0.19	0.1	0.1	0.11	0.17	0.18	0.25	0.15

Note: Headcount is measured as the proportion of population that is deemed poor.

When considering within/between decompositions between age, age-gender, states, caste, and place of residence groups, within group inequality is always much higher than between group inequality. In the individual MPI, the relative contribution of between age-gender groups is higher than in the household MPI (7 vs. 2 percent) while the reverse is the case for states, caste-religion, and place of residence.

Inequality among the poor is actually higher in the household than the individual MPI,[18] contradicting our finding from CSPI analysis. The contradiction in results of the two inequality measures is likely because the positive-multiple of variance measure is an absolute measure, whereas the inequality component in CSPI is a relative measure. Inequality among the poor is also dependent on the value of `k' and might change for different values.

We also calculate Ginis for achievement scores (defined as 1 - deprivation score) for both the individual and the household-based MPI. The household MPI Gini is lower than the one estimated using individual data (0.16 vs. 0.2). The Gini among women is higher than among men for the individual MPI (0.21 vs. 0.18).[19]

Table 4.5: Inequality (variance based measure) across the population and among the poor

	Total Inequality		Gender	Age-Gender	Household	States	Caste & Religion	Place of Residence
			Contribution of within and between components by various groupings (%)					
Inequality across the entire population								
Household MPI	0.171	Within	100	98	0	82	91	80
		Between	0	2	100	18	9	20
Individual MPI	0.228	Within	99	93	30	87	91	84
		Between	1	7	70	13	9	16
Inequality among poor								
Household MPI	0.096	Within	100	98	0	95	98	98
		Between	0	2	100	5	2	2
Individual MPI	0.091	Within	100	98	35	91	97	92
		Between	0	2	65	9	3	8

We also investigate gender differentials in poverty measures across various socio-economic groups (Table 4.6). In all groups females are more likely to be poor and have a higher individual MPI than men. The extent of the differentials between them varies by groups but is always larger for the individual measure than the comparable household MPI. The difference in the adjusted headcount ratio is 0.04 points in metropolitan areas, but doubles to 0.08 points in least developed villages. *Adivasis* (tribals) and *dalits* (lower caste) are among the poorest groups and Christians, Jains and others are the least poor group. Individuals belonging to four member households are the least poor. Single member households are the poorest and most of these are composed of

---

[18] This is not comparing the same group of individuals as those who are identified as poor by the two measures are not completely the same, even though there is a large overlap. See discussion below on overlap.
[19] These results are available on request.

Table 4.6: MPI measure for household and individual indicators by sex and various household and individual characteristics

	Comparable HH MPI			Individual HH MPI		
	Male	Female	Total	Male	Female	Total
Area						
Metro urban	0.02	0.03	0.03	0.05	0.09	0.07
Other urban	0.07	0.08	0.08	0.1	0.15	0.13
More developed villages	0.19	0.2	0.19	0.21	0.28	0.25
Less developed villages	0.28	0.29	0.28	0.28	0.36	0.32
Caste & Religion Groups						
Brahmin	0.11	0.11	0.11	0.09	0.15	0.12
Forward caste	0.09	0.1	0.09	0.1	0.16	0.13
OBC	0.18	0.2	0.19	0.19	0.28	0.24
Dalit	0.24	0.25	0.24	0.26	0.33	0.3
Adivasi	0.29	0.3	0.3	0.3	0.38	0.34
Muslim	0.19	0.2	0.2	0.24	0.28	0.26
Christian, Sikh, Jain	0.03	0.04	0.03	0.05	0.08	0.07
Regions						
North	0.09	0.1	0.09	0.11	0.18	0.14
Central	0.29	0.29	0.29	0.28	0.35	0.32
East	0.25	0.27	0.26	0.26	0.34	0.3
North-East	0.08	0.09	0.09	0.14	0.18	0.16
West	0.12	0.12	0.12	0.13	0.2	0.17
South	0.08	0.09	0.08	0.13	0.19	0.16
Consumption Quintiles						
First	0.34	0.35	0.35	0.35	0.41	0.38
Second	0.24	0.24	0.24	0.26	0.32	0.29
Third	0.18	0.18	0.18	0.2	0.27	0.23
Fourth	0.12	0.12	0.12	0.14	0.21	0.17
Fifth	0.06	0.07	0.06	0.07	0.13	0.1
Household Size						
Single member HH	0.36	0.52	0.48	0.36	0.52	0.48
Two member HH	0.32	0.31	0.32	0.3	0.4	0.35
Three member HH	0.17	0.18	0.17	0.19	0.28	0.23
Four member HH	0.15	0.15	0.15	0.17	0.23	0.2
Five member HH	0.17	0.18	0.17	0.19	0.26	0.22
Six member HH	0.2	0.2	0.2	0.21	0.28	0.25
Seven member HH	0.2	0.21	0.2	0.22	0.29	0.25
Eight or more member HH	0.2	0.2	0.2	0.21	0.27	0.24
Household Head						
Female-Headed HH	0.19	0.22	0.21	0.18	0.3	0.26
Male-Headed HH	0.19	0.19	0.19	0.2	0.27	0.23

Table 4.7: MPI measure for major states in India by gender

	Comparable Household MPI				Individual MPI			
	Male	Female	Total	Rank	Male	Female	Total	Rank
Kerala	0.01	0.02	0.02	1	0.04	0.07	0.05	1
Delhi	0.03	0.04	0.04	2	0.06	0.12	0.09	2
Himachal Pradesh	0.1	0.11	0.1	5	0.1	0.16	0.13	3
Punjab	0.08	0.07	0.07	3	0.11	0.15	0.13	3
Tamil Nadu	0.07	0.08	0.07	3	0.1	0.16	0.13	3
Maharashtra	0.1	0.11	0.11	9	0.12	0.18	0.15	6
Jammu & Kashmir	0.1	0.1	0.1	5	0.13	0.2	0.16	7
Haryana	0.11	0.12	0.12	11	0.13	0.21	0.17	8
Uttarakhand	0.15	0.17	0.16	13	0.14	0.23	0.19	9
Assam	0.1	0.11	0.1	5	0.16	0.21	0.19	9
Karnataka	0.1	0.1	0.1	5	0.16	0.22	0.19	9
Andhra Pradesh	0.11	0.12	0.11	9	0.16	0.24	0.2	12
Gujarat	0.15	0.16	0.15	12	0.17	0.25	0.21	13
West Bengal	0.19	0.2	0.2	14	0.24	0.28	0.26	14
Orissa	0.22	0.24	0.23	15	0.23	0.31	0.27	15
Jharkhand	0.25	0.26	0.25	17	0.25	0.33	0.29	16
Chhattisgarh	0.28	0.28	0.28	18	0.25	0.33	0.29	16
Rajasthan	0.24	0.24	0.24	16	0.25	0.36	0.3	18
Madhya Pradesh	0.29	0.29	0.29	19	0.27	0.35	0.31	19
Uttar Pradesh	0.3	0.3	0.3	20	0.29	0.36	0.33	20
Bihar	0.33	0.36	0.35	21	0.32	0.42	0.37	21

individuals who are widowed or divorced (91%) and are predominantly women (81%). Individuals in female-headed households are more likely to be poor as compared to from male-headed households. The overall rankings of groups do not change much between the household and individual MPI, as the intra-household differences would have to be significantly different for individual measures to change rankings.

Southern, Northern, and North-Eastern states have the lowest household and individual MPI and Central and Eastern states the highest in both. While the individual MPI shows higher levels in all regions, the differential between the regions is smaller in the individual measure than in the household measure which is likely related to high levels of deprivation among older population groups even in states with a low MPI; this shows up in the individual MPI but is masked in the household one. Table 4.7 lists MPI by gender for states for both household and individual measures. The ranking of states are very similar across the two measures with a spearman rank correlation of greater than 0.97. Kerala with an adjusted headcount ratio 0.05 is the best performer and at comparable levels to several middle-income Latin American countries, while Bihar is the worst performing state in India and worse than several sub-Saharan countries like

Nigeria, Tanzania and Sudan.[20]   The state rankings broadly match the rankings obtained by Alkire and Seth (2015) for India for the year 2005-06.

Another way to assess the bias of a household-based measure is to investigate the classifications of individuals into poor and non-poor categories using the household and individual MPI to ascertain the degree of overlap between the two (Table 4.8). 22 percent of men and 27 percent of women are misclassified by the household measure. Men are equally likely to be misclassified as poor or non-poor, while women are more likely to be misclassified as non-poor when using household data. This confirms our hypothesis that the worse off group is more likely to be misclassified as non-poor in the household measure when expansive thresholds are used. Since most disparities are in the expansive education dimension (while there are few gender disparities in the restrictive nutrition category), the education dimension drives the misclassification among women.

In Table 4.9 we decompose MPI to obtain the contribution of each dimension[21]. The health dimension is the biggest contributor to household MPI, while education is the biggest contributor to the individual MPI. This is partly because in the household MPI health is defined in the restrictive fashion leading to higher deprivation rates than with the individual MPI, while the education indicator thresholds are defined in an expansive fashion leading to underestimate of deprivation rates in household as compared to individual MPI. Thus the household-level MPI overemphasizes the health dimension and underplays the education one, explaining the large differences in the decomposition.[22]

These findings are similar to what Haddad and Kanbur (1990) find when estimating uni-dimensional poverty rates based on nutritional status for the population of Philippines. They find that women are more likely to be poor when using individual nutrition data than when using household as the unit of analysis. They also find that intra-household inequality in nutrition is between 30 to 40 percent of total inequality.

---

[20] Based on household MPI calculated for 101 countries in 2015 by Oxford Poverty and Human Development Initiative. Available at http://www.dataforall.org/dashboard/ophi/index.php/mpi/country_briefings
[21] We reweight the indicators within the dimension when data is missing for some indicator due to lack of eligible population in the household. This has the side-effect that we cannot decompose MPI by indicators, but still can do it by dimensions.
[22] In rural areas standard of living and health are almost equal contributors to the household MPI and it plays a bigger role than health in individual MPI. In urban areas standard of living indicators play a smaller role in MPI. This is partly because access to basic services like electricity, sanitation, water and cooking fuel are more readily available in urban areas. But some other aspects of standard of living as density of housing are not captured in the indicators and likely to be worse in urban areas than rural areas. Using 2005-06 NFHS data Alkire and Seth (2015) find standard of living to be highest contributor of household MPI. This seems to have changed with better provision of several public goods like electricity, water, cooking fuel etc. and increased standard of living during the high economic growth period between 2006 and 2012.

Table 4.8: Classification of individuals by household and individual MPI

	Individual MPI		
Comparable HH MPI	Non-Poor	Poor	Total
Male			
Non-Poor	0.51	0.11	0.62
Poor	0.11	0.27	0.38
Female			
Non-Poor	0.42	0.19	0.61
Poor	0.08	0.31	0.39
Total			
Non-Poor	0.47	0.15	0.62
Poor	0.09	0.29	0.38

4.5.1. Intrahousehold measures

The measures investigated so far are all sample level measures. We have also used some intra-household measures to better understand the differences in deprivation levels between men and women in the same household. Based on Gender Parity Index (GPI) defined by Alkire et. al (2013), we have developed measures to determine level of disparity in deprivation scores among different groups of men and women in the household. Alkire et. al (2013) define GPI as a relative inequality measure that measures inequality in empowerment between the primary adult male and female in each household. We extend this further to look at inequality in deprivation scores between the primary couple as well as other groups of males and females (all adults, all children) in the household.

Table 4.9: Contribution of each dimension to MPI measure (%) (Dual cutoff method with k=0.33)

	Rural			Urban			All		
	Male	Female	Total	Male	Female	Total	Male	Female	Total
Household MPI									
Education	17	16	17	20	29	17	17	21	17
Health	42	44	42	60	43	50	44	42	44
Standard of Living	42	40	42	20	29	33	39	37	39
Individual MPI									
Education	46	50	47	67	64	64	50	52	52
Health	25	21	23	22	21	18	25	22	22
Standard of Living	29	29	30	11	14	18	25	26	26

For GPI, we construct a new censored deprivation score where all individuals whose deprivation score is below the poverty threshold of k, get a deprivation score equal to k. So new censored deprivation score $c_i' = c_i$ if $c_i \geq k$ and $c_i' = k$ if $c_i < k$. Each household is classified based on gender parity. If the average deprivation score of the group of women being considered for

the measure in the household is above the poverty cutoff (k) and the women's average deprivation score is higher than their male counterpart then the household is deemed to lack gender parity. A household has parity if either the average women is non-poor (deprivation score below poverty cutoff) or women's deprivation score is lower than men's.

Following Alkire et. al (2013), GPI can be broken up into headcount component and an intensity component. The headcount gives us the proportion of households that do not have gender parity and the intensity component gives the average deprivation gap between women and men's group in the households that lack gender parity.

$$H_{GPI} = h/m$$

$$I_{GPI} = \frac{1}{h}\sum_{j=1}^{h}\frac{c_j'(k)^W - c_j'(k)^M}{1 - c_j'(k)^M}$$

where h is the proportion of households that lack gender parity and $c_j'(k)^W$ and $c_j'(k)^M$ are the censored deprivation scores for the group of women and men living in household j.

GPI is then defined as

$$GPI = 1 - (H_{GPI} * I_{GPI})$$

GPI can be improved either by reducing the proportion of households that lack gender parity or by decreasing the average difference in deprivation scores in households that lack gender parity.

Table 4.10: Intra-household inequality measures

	% of HH's with higher average women's deprivation score than men's.	% of HH's with higher average men's deprivation score than women's.	Average difference in deprivation score	% HH's with no Gender Parity	Average deprivation gap	Gender Parity Index
Primary Couple in the household	54	11	0.075	45	0.23	0.90
All adult members of household	63	15	0.08	44	0.23	0.90
All children in the household	19	18	-0.004	14	0.24	0.97
All household members	61	21	0.06	38	0.21	0.92

In addition to calculating the gender parity index (GPI) that considers only a subset of households that lack parity, we also investigate differences in deprivation scores within the household for different groups of men and women. Table 4.10 reports the intra-household

measures for four groups of men and women: differences between the primary couple of the household, between all adult men and adult women in the household, between girls and boys in the household and between all males and females in the household. 61 percent of households have average female deprivation score that is higher than average male deprivation score, underlying the low status of women in the households. When we look at only children in the household there does not appear to be a difference in deprivation scores between girls and boys. The Gender Parity Index is highest among children, and lowest when only adults or the primary couple is included in calculating the index. These findings are similar to sample level analysis that pointed to increase in differences in poverty rates by gender as age increases. Table 4.11 reports the intra-household measures for various socio-demographic groups and these results mirror those of sample level analysis.

Table 4.11: Intra-household inequality measures for various socio-economic groupings

	% of HH's with higher average women's deprivation score than men's.	% of HH's with higher average men's deprivation score than women's.	Average difference in deprivation score	% HH's with no Gender Parity	Average deprivation gap	Gender Parity Index
Area						
Rural	69	23	0.07	48	0.22	0.89
Urban	44	18	0.04	18	0.15	0.97
Region						
North	59	18	0.06	27	0.17	0.95
Central	68	22	0.07	51	0.24	0.88
East	67	23	0.06	48	0.21	0.90
North-East	49	21	0.04	25	0.18	0.96
West	55	18	0.06	30	0.2	0.94
South	54	21	0.05	28	0.17	0.95
Caste-Religion Groups						
SC Households	67	23	0.06	46	0.22	0.90
ST Households	69	22	0.07	57	0.22	0.87
Muslim	60	28	0.04	36	0.18	0.94
OBC	63	21	0.07	40	0.21	0.92
Others	50	16	0.05	22	0.2	0.96
Consumption Quintiles						
First	69	28	0.05	57	0.21	0.88
Second	70	24	0.06	48	0.21	0.90
Third	66	24	0.06	41	0.22	0.91
Fourth	60	19	0.06	32	0.21	0.93
Fifth	45	13	0.05	22	0.18	0.96

4.5.2. Determinants of MPI

To delve into the multivariate correlates of the individual versus the household MPI we run a regression of deprivation scores on a range of individual and household characteristics with various fixed effects (Table 4.12 and 4.13). The regression model can be expressed as:

$$c_i = \alpha + \beta_1(Female) + \beta_2 Age + \beta_3 Age^2 + \beta_4 Age^3 + \gamma X + \varepsilon$$

where X is a set of individual, household controls and fixed effects (state, PSU or HH).

In all our specifications we find that females have a significantly higher deprivation score than males even after controlling for various other parameters. Even with household fixed effects, which eliminates all inter-household variation in estimation and controlling for other individual characteristics like age, occupation, marital status and relationship to head we find females to be

Table 4.12: Determinants of Household MPI deprivation score

	(1)	(2)	(3)	(4)
	OLS	State FE	State FE with Interactions	PSU FE
	b/se	b/se	b/se	b/se
Female	0.004***	0.007***	0.119***	0.007***
	0.001	0.002	0.016	0.001
Male Headed HH	0.030***	0.029***	0.028***	0.026***
	0.001	0.005	0.005	0.003
Age	-0.007***	-0.006***	-0.006***	-0.004***
	0.000	0.000	0.000	0.000
Age # Age	$.007*10^{-2}$***	$0.007*10^{-2}$***	$0.007*10^{-2}$***	$0.006*10^{-2}$***
	0.000	0.000	0.000	0.000
Age # Age # Age	$-0.002*10^{-4}$***	$-0.002*10^{-4}$***	$-0.002*10^{-4}$***	$-0.002*10^{-4}$***
	0.000	0.000	0.000	0.000
Marital Status (base:Married)				
Married, spouse absent	0.013***	-0.003	-0.001	-0.009***
	0.003	0.003	0.003	0.003
Unmarried 2	-0.002	0.003	0.004	0.006***
	0.002	0.005	0.005	0.002
Widowed 3	0.002	-0.001	-0.001	-0.004*
	0.002	0.003	0.003	0.002
Separated/Divorced 4	0.035***	0.034***	0.036***	0.029***
	0.007	0.008	0.008	0.006
Married no gauna 5	0.064***	0.036***	0.037***	0.010
	0.010	0.009	0.009	0.010
Relationship to Head (base: Self)				
Wife/Husband 2	-0.008***	-0.002	-0.004	-0.003***
	0.002	0.002	0.002	0.001
Son/Daughter 3	-0.058***	-0.043***	-0.044***	-0.029***
	0.002	0.004	0.004	0.002
Child-in-Law 4	-0.080***	-0.058***	-0.059***	-0.042***
	0.002	0.005	0.005	0.003
Grandchild 5	-0.140***	-0.110***	-0.111***	-0.077***
	0.003	0.009	0.009	0.004

Father/Mother 6	0.014***	0.019***	0.018***	0.012***
	0.003	0.003	0.004	0.003
Other	-0.055***	-0.035***	-0.036***	-0.022***
	0.002	0.007	0.007	0.003
Caste & Religion (base:OBC)				
Brahmin 1	-0.016***	-0.024***	-0.022***	-0.026***
	0.002	0.006	0.006	0.004
Forward caste 2	-0.041***	-0.028***	-0.027***	-0.023***
	0.001	0.004	0.004	0.003
Dalit 4	0.018***	0.027***	0.026***	0.035***
	0.001	0.004	0.004	0.003
Adivasi 5	0.032***	0.055***	0.053***	0.044***
	0.002	0.008	0.008	0.006
Muslim 6	-0.005***	0.009	0.007	0.007
	0.001	0.007	0.007	0.005
Christian, Sikh, Jain 7	-0.089***	-0.041***	-0.043***	-0.029***
	0.002	0.008	0.010	0.006
HH-Size (base: One-Member HH)				
Two member HH	-0.115***	-0.113***	-0.017	-0.117***
	0.007	0.010	0.014	0.007
Three member HH	-0.208***	-0.202***	-0.107***	-0.202***
	0.007	0.013	0.015	0.007
Four member HH	-0.221***	-0.214***	-0.117***	-0.212***
	0.007	0.013	0.015	0.007
Five member HH	-0.209***	-0.209***	-0.115***	-0.211***
	0.007	0.013	0.014	0.007
Six member HH	-0.200***	-0.207***	-0.111***	-0.209***
	0.007	0.014	0.016	0.007
Seven member HH	-0.198***	-0.209***	-0.114***	-0.215***
	0.007	0.013	0.016	0.007
Eight or More member HH	-0.197***	-0.214***	-0.118***	-0.222***
	0.007	0.014	0.015	0.007
Place of Residence (base: Metropolitan)				
other urban 1	0.047***	0.036***	0.033***	
	0.001	0.008	0.008	
more dev vill 2	0.109***	0.113***	0.108***	
	0.001	0.010	0.010	
less dev vill 3	0.178***	0.146***	0.139***	
	0.001	0.009	0.009	
Female # Age			0.0004	
			0.000	
Female # Age # Age			$-0.001*10^{-2}$	
			0.000	
Female # Age # Age # Age			$0.0009*10^{-4}$	
			0.000	
Female # Male Headed HH			0.005***	
			0.002	
Constant	0.590***	0.549***	0.463***	0.587***
	0.008	0.017	0.020	0.008
Education of HH-Head	Yes	Yes	Yes	Yes
Occupation Status	Yes	Yes	Yes	Yes
Sex # HH-Size	No	No	Yes	No
Sex # Caste & Religion Groups	No	No	Yes	No

Sex # Place of Residence	No	No	Yes	No
State # Sex	No	No	Yes	No
R2-Within	0.353	0.319	0.320	0.163
R2-Between		.423	.456	.484
R2-Overall		.341	.342	.254
N	198614	198614	198614	198614
Number of Groups		33	33	2462

Dependent variable is deprivation score for Comparative Household MPI. Robust Standard Errors are reported
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

worse off. But this effect in the household MPI deprivation score is driven entirely by differences in household composition and not directly due to gender disparities within households. The gender effect is also substantially higher for the individual MPI than for household MPI, even when controlling for other covariates and fixed effects.

In the absence of individual analysis most gendered poverty analysis have relied on comparison of sex of household head to proxy for gender. This is a flawed approach, since even male-headed households have female members who might have different level of deprivations than the male head. Several studies have claimed that female-headed households are more likely to be poor than male-headed households and hence females are poorer than males (e.g. Dreze and Srinivasan, 1997; Chant 2004; Klasen, Lechtenfeld and Povel 2015). Our regression analysis finds that after controlling for other factors, most importantly the education of head of household, female-headed households on average have *lower* deprivation scores than males. This finding holds true for both individual and household MPI measures. Thus while adult women *are* disadvantaged in poverty in India, as demonstrated by our individual MPI, this has nothing to do with household headship, but is an intra-household inequality issue.

Figure 4.1: Differential Impact of Sex of the Individual by Age
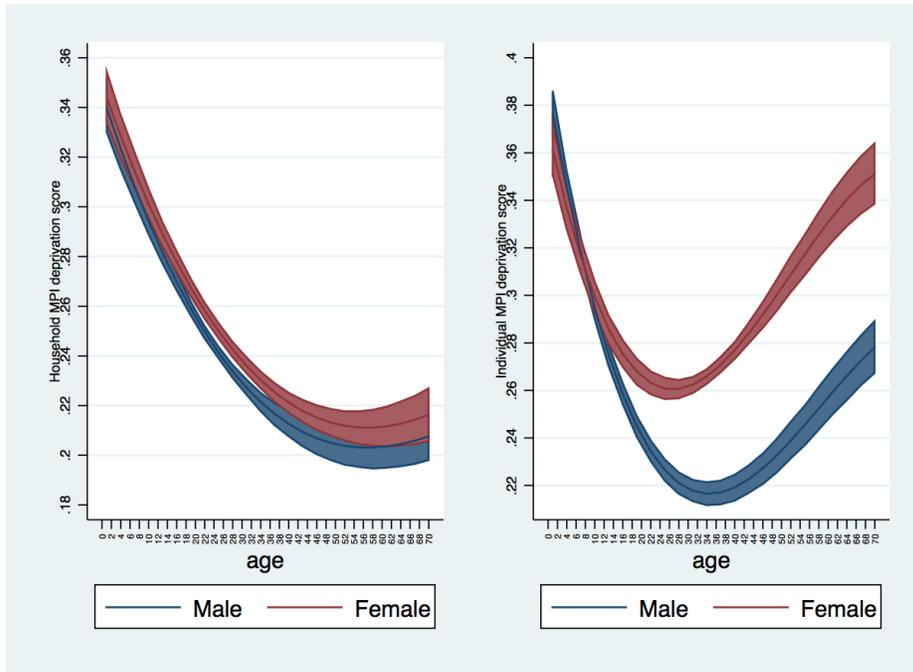
Age of the individual has a stronger impact on deprivation scores in the individual MPI measure. The combined partial effect of age and gender on the household versus individual MPI is substantial as shown in Figure 4.1. While in the household-based measure, the deprivation scores fall uniformly with age for

Table 4.13: Determinants of Individual MPI deprivation score

	(1)	(2)	(3)	(4)	(5)
	OLS	State FE	State FE with Interactions	PSU FE	HH FE
	b/se	b/se	b/se	b/se	b/se
Female	0.014***	0.017***	-0.013	0.017***	0.014***
	0.001	0.003	0.018	0.001	0.001
Male Headed HH	0.035***	0.034***	0.039***	0.028***	
	0.002	0.005	0.006	0.003	
Age	-0.012***	-0.011***	-0.012***	-0.009***	-0.007***
	0.000	0.000	0.001	0.000	0.000
Age # Age	$0.025*10^{-2}$***	$0.025*10^{-2}$***	$0.024*10^{-2}$***	$0.024*10^{-2}$***	$0.022*10^{-2}$***
	0.000	0.000	0.000	0.000	0.000
Age # Age # Age	$-0.015*10^{-4}$***	$-0.015*10^{-4}$***	$-0.013*10^{-4}$***	$-0.016*10^{-4}$***	$-0.016*10^{-4}$***
	0.000	0.000	0.000	0.000	0.000
Marital Status (base: Married)					
Married, spouse absent	0.043***	0.028***	0.012**	0.024***	0.031***
	0.003	0.005	0.005	0.004	0.003
Unmarried	0.023***	0.027***	0.016**	0.030***	0.016***
	0.002	0.007	0.007	0.002	0.002

Widowed	0.064***	0.060***	0.036***	0.057***	0.056***
	0.002	0.003	0.003	0.003	0.002
Separated/Divorced	0.085***	0.085***	0.070***	0.082***	0.054***
	0.007	0.006	0.006	0.007	0.006
Married no gauna	0.063***	0.029***	0.017**	0.010	-0.010
	0.011	0.007	0.007	0.013	0.009
Relationship to Head (base: Self)					
Wife/Husband	0.066***	0.072***	0.035***	0.071***	0.074***
	0.002	0.006	0.006	0.002	0.001
Son/Daughter	-0.087***	-0.070***	-0.078***	-0.056***	-0.015***
	0.002	0.003	0.003	0.002	0.002
Child-in-Law	-0.074***	-0.051***	-0.076***	-0.034***	0.019***
	0.003	0.005	0.005	0.003	0.003
Grandchild	-0.171***	-0.139***	-0.146***	-0.104***	-0.011***
	0.003	0.005	0.005	0.004	0.004
Father/Mother	0.090***	0.094***	0.076***	0.086***	0.056***
	0.003	0.005	0.005	0.003	0.003
Other	-0.045***	-0.025***	-0.035***	-0.011***	0.025***
	0.003	0.007	0.006	0.003	0.002
Caste & Religion (base: OBC)					
Brahmin	-0.027***	-0.037***	-0.027***	-0.037***	
	0.002	0.005	0.005	0.004	
Forward caste	-0.045***	-0.033***	-0.027***	-0.029***	
	0.001	0.006	0.006	0.003	
Dalit	0.021***	0.029***	0.028***	0.034***	
	0.001	0.003	0.003	0.003	
Adivasi	0.033***	0.054***	0.049***	0.044***	
	0.002	0.007	0.007	0.005	
Muslim	0.006***	0.016**	0.019**	0.011***	
	0.001	0.008	0.008	0.004	
Christian, Sikh, Jain	-0.093***	-0.042***	-0.032***	-0.031***	
	0.002	0.005	0.006	0.005	
HH-Size (base: One-Member HH)					
Two member HH	-0.044***	-0.044***	-0.015	-0.047***	
	0.006	0.006	0.013	0.006	
Three mem. HH	-0.071***	-0.066***	-0.034***	-0.065***	
	0.006	0.007	0.012	0.006	
Four member HH	-0.077***	-0.070***	-0.038***	-0.068***	
	0.006	0.007	0.012	0.006	
Five member HH	-0.066***	-0.067***	-0.039***	-0.068***	
	0.006	0.006	0.012	0.006	
Six member HH	-0.056***	-0.065***	-0.038***	-0.067***	
	0.006	0.007	0.014	0.006	
Seven mem. HH	-0.053***	-0.067***	-0.040***	-0.072***	
	0.006	0.008	0.014	0.006	
Eight or More member HH	-0.055***	-0.075***	-0.049***	-0.083***	
	0.006	0.007	0.013	0.006	
Place of Residence (base: Metropolitan)					
other urban	0.033***	0.026***	0.020***		
	0.001	0.007	0.007		

more dev vill	0.088***	0.095***	0.075***		
	0.002	0.008	0.008		
less dev vill	0.158***	0.128***	0.104***		
	0.002	0.007	0.008		
Female # Age			0.002***		
			0.000		
Female # Age # Age			$-0.0001*10^{-2}$		
			0.000		
Female # Age # Age # Age			$-0.002*10^{-4}$*		
			0.000		
Female # Male Headed HH			0.001		
			0.003		
Constant	0.514***	0.473***	0.499***	0.495***	0.272***
	0.007	0.015	0.023	0.008	0.004
Ed. HH-Head	Yes	Yes	Yes	Yes	No
Occ. Status	Yes	Yes	Yes	Yes	Yes
Sex # HH-Size	No	No	Yes	No	No
Sex # Caste & Religion Groups	No	No	Yes	No	No
Sex # Residence	No	No	Yes	No	No
State # Sex	No	No	Yes	No	No
R2-Within	0.433	0.417	0.421	0.307	0.260
R2-Between		.548	.667	.573	.0617
R2-Overall		.423	.437	.36	.0941
N	197942	197942	197942	197942	198236
N Groups		33	33	2462	40726

Dependent variable is deprivation score for Individual MPI measure. Robust Standard Errors are reported

- p<0.10, ** p<0.05, *** p<0.01

males and females and there is little gender difference, there is a pronounced U-shape in age for both genders in the individual measure, with females being much more deprived in the individual measure.[23]

In addition, widowed women's deprivation score in the household measure is not significantly different from that of married women for most specifications, but in the individual measure they have significantly higher deprivation scores. Married women are significantly better off in the individual measure than all other women, while that's not always the case in the household-based measure. The wife or husband of the head of the household is always significantly worse-off than the head in individual-based deprivation score, while they are better-off in the household-based measure. Muslims are significantly worse-off than OBC's when comparing

---

[23] Note that the high deprivation among young children in both measures is related to the imputation of education scores of adults to children below 6 and should therefore be treated with caution.

their individual deprivation scores, but this does not hold in the household-based measure. Single-member households are worst-off than bigger households and metropolitan areas are best-off in both the individual and household measures.[24]

Besides the impact of various individual and household characteristics, the fixed effects regression provide useful information on state-specific differences in deprivation scores after controlling for other factors.

Figure 4.2 shows the estimate of fixed effects for both household and individual MPI measures from these regressions. The pattern across states is similar for both individual and household MPI measures, confirming our earlier finding based on summary statistics that ranking of states does not change substantially when using either of the measures. Bihar, Uttar Pradesh and Madhya Pradesh have the highest deprivation scores after controlling for other parameters and north-eastern on average have the lowest deprivation scores. Even though state per-capita GDP and economic growth rates for north-eastern states are lower than in western and southern states, they perform better in multi-dimensional poverty measures. In sum, these regressions indicate that the individual MPI measure leads us to substantially different conclusions on deprivation by age, gender, marital status, relationship to the head of the household, and caste/religion groups.

4.6.   Robustness Analysis

Designing a poverty measure involves a selection of various parameters, and we are interested in determining how sensitive our major results of gender differential in poverty are to these parameter choices. We investigate the robustness of our results to i) change in deprivation thresholds for education dimension for children under seven years of age ii) change in how health deprivation status of individuals with missing data on nutrition is determined ii) weights assigned to the three dimensions (w) and iii) poverty cutoff (k).

Figure 4.2: Regression results for state fixed effects for Household and Individual MPI

---

[24] We also analyze differential impact of gender across various groups. Gender differences in deprivation scores are significant only among adults; they are absent in urban areas and substantial in the rural population; absent in northeastern and several southern states and large among central and eastern states. Detailed results are available upon request.
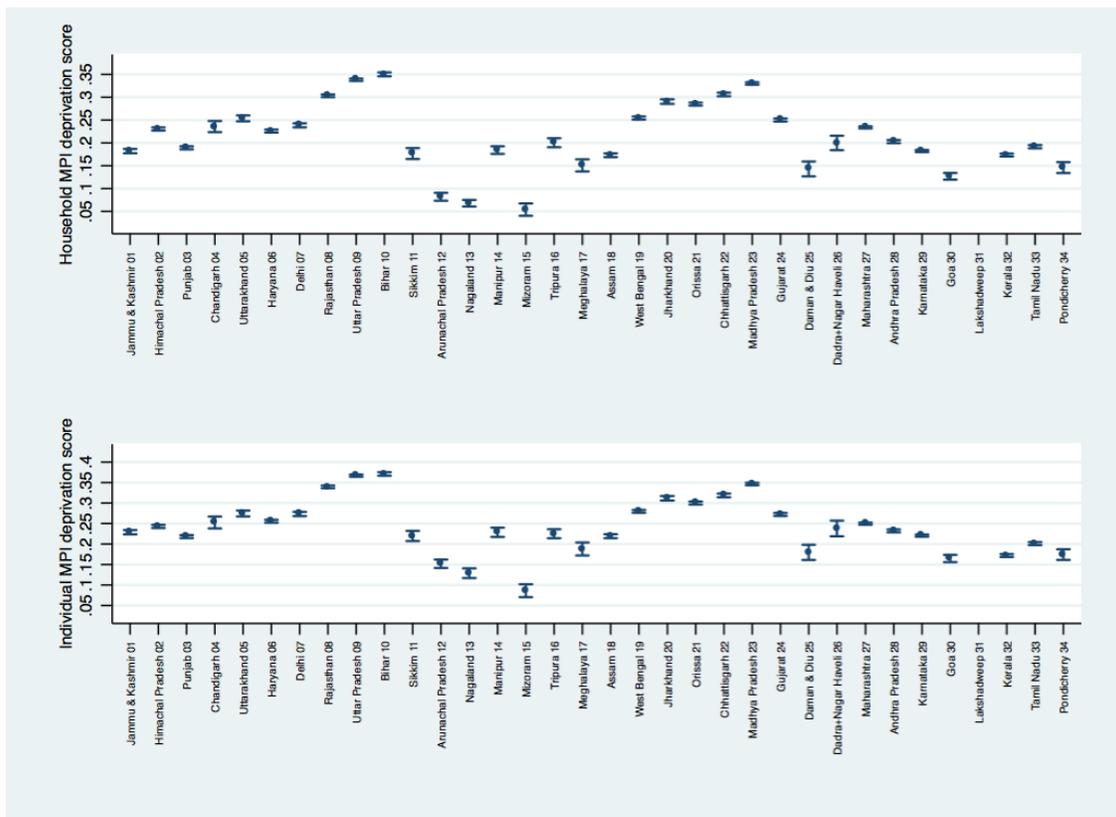
Table 4.14 presents the poverty headcount, intensity and MPI for five different individual MPI measures constructed by changing the parameters. In our benchmark individual measure children under seven years of age are deemed deprived in schooling if half or more members in the household above 12 years of age have not completed five years of education. We do not have any alternate information on education potential of these children and chose to define the deprivation based on other household members. But since access to primary schools is expanding rapidly the likelihood of these children completing five years of education is higher than adults in the household. One alternative assumption we could make in defining deprivation threshold for children below seven would be to assume that they are non-deprived. We do this in our first alternate measure and find that even though the level of MPI decreases from our benchmark individual measure the differential between men and women still exists. The absolute differential

126

between men and women in MPI remains the same and the relative differential increases slightly.[25]

Next, we consider if our conclusions are robust to how health deprivation status of individuals with missing nutrition information is determined. In IHDS survey, we have individual data on nutrition for 65 percent of all individuals and infer deprivation status for the remaining 35 percent of the population based on the status of the demographic group they belong to (as explained previously). We are missing data on 26 percent of females and 44 percent of males. The missing data is concentrated more among adult males (54 percent). In the benchmark version, the deprivation status of these adult males in the household would be determined based on information on nutrition status of other adult males for whom data was collected (deprived if 50 percent of adults males are underweight). If data was not collected on any adult males in the household then their deprivation status is determined by status of other adults in the household similarly. In this we are making the assumption that deprivation status of individuals with missing data can be implied based on the status of the broader demographic group they belong to. Instead of this we test if assigning individuals with missing data the same status as they would have been assigned when constructing household MPI makes a difference to our results. For individuals with missing data we deem them as deprived in nutrition if any one individual in the household is under-nourished just as in Household MPI. But this does not change our results qualitatively; women are still more likely to be deemed as poor than males.

Next, we ask if our conclusions are robust to a range of weights. To test this, we estimated individual MPI using three additional weighting schemes: i) giving 50 percent to education and 25 percent each to health and standard of living, ii) giving 50 percent to health and 25 percent each to education and standard of living, to equalize the weight of the expansive education indicator and the restrictive nutrition one to 25% each and iii) giving 50 percent to standard of living and 25 percent each to education and health. Within each dimension all indicators got equal weights, except aspects that had time use indicators[26]. We find that the levels of MPI and the extent of difference between deprivation scores of men and women changes for the different weighting schemes, but women are significantly worse off than men in all the three alternative weighting schemes; also poverty rates are larger than in the household measure, demonstrating that this is not only related to weights. The differences are larger and overall poverty is higher

---

[25] Using the benchmark measure we find a small gender differential in children below seven years of age. 54 percent of boys and 56 percent of girls below seven years of age are deemed poor in the benchmark measure (Table 4.3). The corresponding numbers when using the no education deprivation assumption are 18 percent and 19 percent respectively.
[26] In case of water and cooking the access indicator and time use indicator got half of the weight assigned to each standard of living indicator.

when education is given higher weight as the differences in schooling between men and women are large.

Table 4.14: Robustness Analysis

	Men	Women	Total	Difference between women and men's measure	
				Absolute	Relative
Individual MPI					
Headcount	0.38	0.5	0.44	0.12	32%
Intensity	0.53	0.54	0.54	0.01	2%
MPI	0.2	0.27	0.24	0.07	35%
Individual Alternate MPI (Children under 7 years are assumed to be non-deprived in education)					
Headcount	0.33	0.46	0.4	0.13	39%
Intensity	0.52	0.53	0.53	0.01	2%
MPI	0.17	0.24	0.21	0.07	41%
Individual Alternate MPI (For individuals with missing nutrition information using household nutrition status)					
Headcount	0.4	0.51	0.45	0.11	28%
Intensity	0.54	0.55	0.54	0.01	2%
MPI	0.22	0.28	0.24	0.06	30%
Individual Alternate Weights MPI (Education 0.5, Health 0.25 and Standard of Living 0.25)					
Headcount	0.32	0.45	0.38	0.13	41%
Intensity	0.65	0.66	0.65	0.01	2%
MPI	0.21	0.3	0.25	0.09	43%
Individual Alternate Weight MPI (Education 0.25, Health 0.5 and Standard of Living 0.25)					
Headcount	0.36	0.44	0.4	0.08	22%
Intensity	0.51	0.52	0.52	0.01	2%
MPI	0.18	0.23	0.21	0.05	28%
Individual Alternate Weights MPI (Education 0.25, Health 0.25 and Standard of Living 0.5)					
Headcount	0.4	0.49	0.45	0.09	23%
Intensity	0.51	0.53	0.52	0.02	4%
MPI	0.2	0.26	0.23	0.06	30%

Finally, we investigate if our findings are robust to changes in the poverty cutoff (k). To do that we use the complementary cumulative distribution function (CCDF) – the complement of a cumulative distribution function introduced for this analysis by Alkire et al (2015). The CCDF tells us the proportion of population above any value b and helps us determine the proportion of the population who will be deemed poor if the poverty cutoff is set to b i.e. k=b. Alkire et al (2015) show that if we find first order stochastic dominance between CCDF's for two distributions c and c', then we can claim that distribution c has no lower multidimensional headcount ratio H and adjusted headcount ratio than distribution c' for all values

128

Figure 4.3: Multidimensional poverty headcount for various values of deprivation score cutoff by gender
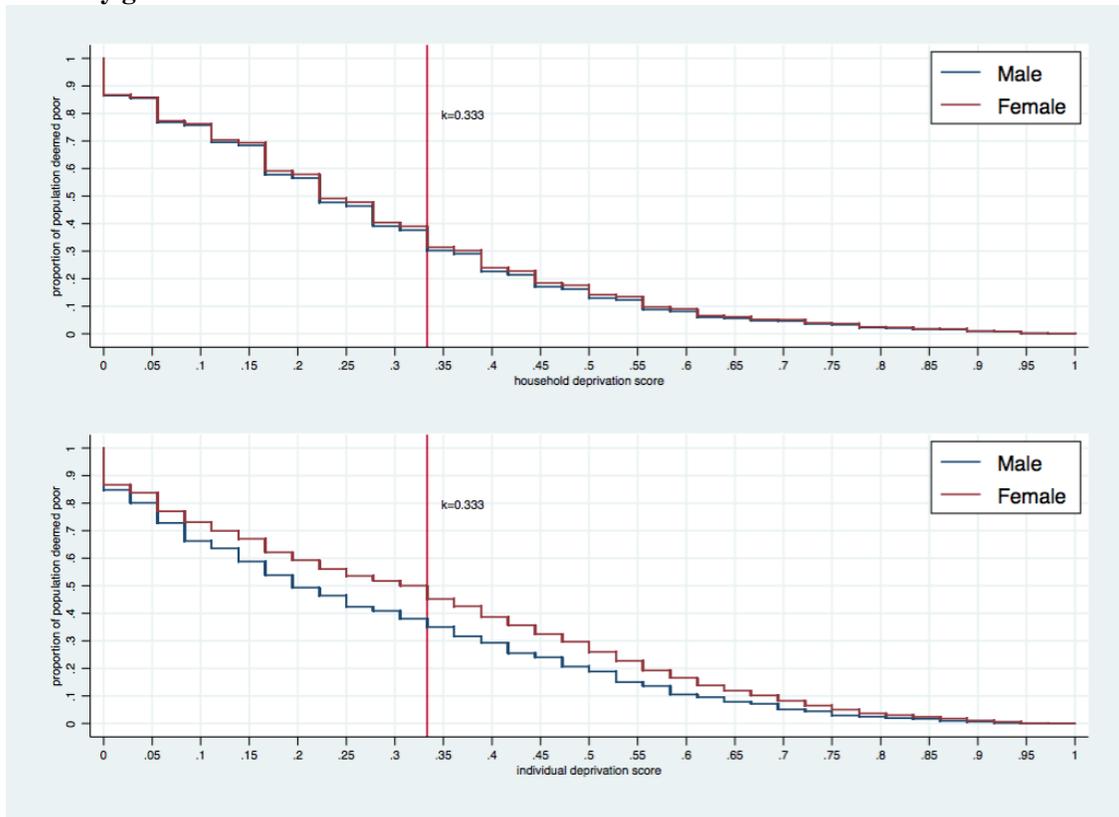
of k. Figure 4.3 plots the CCDFs for men and women for various values of k and we find that the distribution for women dominates that of men. In other words, women's headcount ratio and adjusted headcount are not lower than men's for all values of k. For values of k between 0.1 and 0.8 women have a higher poverty headcount than men. We should also note that the differential is particularly pronounced around the cut-off of 0.33 which we chose, following the practice of the MPI. It is slightly smaller at higher and lower cut-offs, and becomes very small as we move get close to the union (k=1) or the intersection (k=0) approach.

4.7.   Conclusion

In this paper we contribute to the literature on multidimensional poverty measurement by proposing and applying an individual multidimensional poverty measure for India. We find that existing multidimensional poverty measures use household-based assessments for multidimensional poverty measurement, even though individual achievement data are available

for some dimensions of well-being. The use of household-based thresholds based on individual achievement data lead to biases in multidimensional poverty assessment. In the Indian case, we find that household-based MPIs substantially understate poverty, gender inequality, inequality in deprivation across the population, and differentials by age groups. Such misclassification could also affect assessments of poverty trends and targeting. While targeting based on regions or groups other than age or gender would not be very seriously biased when using a household-based measure, targeting based on gender and age groups would. And using the incidence of female-headship as a sign of gendered poverty would be deeply misleading.

Our analysis can only be seen as a first step in this direction. We are only able to individualize deprivation data in some dimensions where available data allow such disaggregation. Following our findings from India, we are therefore likely to understate inequalities in deprivation, particularly in a developing country context. More data would be required, for example, the individualize deprivation in morbidity as well as possession and use of assets. Moreover, our assessment relies on some assumptions about group-based deprivations that are required to create individual deprivation measures for everyone in the household. Clearly here, alternative approaches (such as assessment of individual deprivation by groups) are a possible alternative, as are different assumptions to create deprivation scores for everyone.

But we have demonstrated that the neglect of intrahousehold inequality is a serious issue and actually underestimates in the Indian case poverty and inequality in deprivation by some 30 percent.

5. Are Educated Leaders Good for Education?
Evidence from India

Abstract

Formal education level of leaders has been used as a proxy for quality of leaders. Recently candidates with low education levels have been disqualified from contesting local elections in India. But there is no conclusive evidence linking education to effectiveness of leaders, especially in the Indian context. In this paper, we investigate the linkage between education of the leader and competence by analyzing if educated political representatives result in better education outcomes for children. Using comprehensive data on education in India from multiple data-sources (Annual Status of Education Report, National Sample Survey data, and annual census of schools) and education outcomes (learning levels, enrollment, school funding and infrastructure), we find no evidence that educated political leaders are more competent at delivering better education outcomes for their constituents than less educated leaders. The null effect is precisely estimated. The identification strategy is based on a fuzzy regression discontinuity design (RDD) that exploits quasi-experimental election outcomes of close elections between college graduate and non-graduate leaders. We extensively test our identification setup and fuzzy RDD assumptions. We find that our results are robust to the various assumption in baseline model, different specifications and across various groupings of individuals, households and districts[1].

> Those who insist on literacy as a test and insist upon making it a condition precedent to enfranchisement in my opinion, commit two mistakes. Their first mistake consists in their belief that an illiterate person is necessarily an unintelligent person... Their second mistake lies in supposing that literacy necessarily imports a higher level of intelligence or knowledge than what the illiterate possesses.

> *B.R.Ambedkar, Chief Architect of the Indian Constitution, 1928 to Simon Commission.*

> It is only education which gives a human being the power to discriminate between right and wrong, good and bad.

> *Supreme Court of India, 2015 upholding the law mandating minimum education level for candidates in local elections.*

5.1 Introduction

Two states in India - Rajasthan and Haryana - enacted ballot access restrictions for local body elections in 2014-2015. Both states mandated minimum education requirements for contesting local body elections[2]. These requirements disqualify more than 50 percent of women, 68 percent of Scheduled Caste women and 41 percent of Scheduled Caste men from contesting elections in Haryana (Bhaskar, 2016). This also led to higher proportion of elections being contested by one or no candidates and lower electoral competition. These laws were challenged in the Supreme Court of India which upheld the laws. The Court saw education as a precondition for efficiency and honesty arguing that education will "enable the candidates to effectively discharge duties of the panchayat".[3]

The education requirement and the Supreme Court judgment have been controversial. Many have argued in articles and debates that the law and decision are discriminatory, retrograde, disenfranchising, elitist, unconstitutional and undemocratic (Baxi, 2015; Abdul, 2015; Jaffrelot, 2016; NDTV, 2016). The argument for imposing education requirements on candidates contesting elections rests on the premise that formal education makes leaders competent, honest and accountable[4]. Formal education is seen as a desirable characteristic which is argued to increase the quality of politicians. Educated politicians are seen as more effective at designing and implementing policy, understanding concerns of citizens, dealing with complex policy issues and are considered more accountable.[5] We do not test the impact of this particular policy in this paper, but it serves as one of the motivations to analyze linkages between formal education of the leader and competency.[6]

The existing literature has often considered education to be a proxy for the quality of a leader (Atkinson et al., 2016; Besley et al., 2005). However, there are only a few studies that specifically analyze the impact of formal education of the leader on outcomes, and they give conflicting results. Impact of education of the leader on development outcomes has not been studied systematically

---

[2]In Haryana the education requirements were - class X (10 years of schooling) for general candidates, Class VIII (8 years of schooling) for women and Scheduled Caste men and Class V (5 years of schooling)for Scheduled Caste women.

[3]The Supreme court judgment follows its earlier ruling in 2003 that upheld the requirement that the citizens who have more than two children cannot contest elections.

[4]These arguments were explicitly used as justification by the Rajasthan and Haryana government when introduction the education mandate. https://indconlawphil.wordpress.com/2015/10/01/election-disqualifications-and-the-constituent-assembly-debates/

[5]The constituent assembly of India debated this topic extensively which is discussed here: https://indconlawphil.wordpress.com/2015/10/01/election-disqualifications-and-the-constituent-assembly-debates/.

[6]We do not yet have data on education of local leaders and the policy is too recent to test impact, but we plan to do so in the future.

in India, even though various other characteristics of the leader such as gender, religion, caste and criminality have been extensively studied. Using data on education of national leaders in a cross-country database and random transitions a few studies have shown that educated leaders tend to increase economic growth, foreign investment and education attainment of citizens (Besley, Montalvo, and Reynal-Querol, 2011; Diaz-Serrano and Pérez, 2013; Congleton and Zhang, 2013). These studies only analyze national leaders across countries where disparities across nations are difficult to account for and also they do not analyze impact of state or local leaders' education. But Carnes and Lupu (2016) investigating impact of education of political leaders at national, state and local levels in different contexts find educated politicians to perform no better or worse than less-educated leaders across a range of outcomes.

The aim of the paper is to test the hypothesis that educated politicians are more competent and deliver better outcomes for their constituents. We investigate whether state legislative representatives in India with college degree are better than leaders without college degree for school education outcomes in the districts from which they are elected. If leaders with more formal education are competent and utilize this for the benefit of their constituents then one might expect better education outcomes as compared to having a less educated leaders. Data on various schooling outcomes from multiple data sources are combined with education information on state legislatures who could influence education policy to analyze the impact of educated leaders. We use data on learning outcomes for rural areas from the Annual Status of Education Report (ASER) from 2006 to 2014, enrollment outcome data from various National Sample Survey (NSS) rounds from 2007 to 2014, and annual data on school infrastructure and funding from census of schools (DISE) from 2010 to 2014. Along with these outcomes, we use data on leaders' education level from Association of Democratic Reform (ADR) from 2004 to 2014. The merged dataset allows us to examine effects of political leaders' education on education outcomes for 10 years, the entire span of time for which data on formal education of leaders are available. Our empirical analysis that uses high-quality data with a large sample size enables us to hold a range of institutional and cultural factors constant, and yields statistical power to detect even small effects on education outcomes of leaders.

The main identification challenge is that education level of the leader could be correlated with voter preferences and hence endogenous. To identify the causal impact of educated politicians, we use a fuzzy regression discontinuity design where the proportion of leaders who have completed college is instrumented with the proportion of college graduate leaders who won in close elections against a non-college graduate leader. The validity of the identification strategy relies on the assumption of quasi-randomness of the outcome of a close election (Eggers et al., 2015) – that the preferences

of voters who elect an educated politician in a close election can be assumed to be the same as the preferences of voters who elect a non-college graduate leader. We test these assumptions extensively and show that results of close elections cannot be predicted on any observable characteristics and are quasi-random. This strategy has been used extensively in the literature to study the impact of other characteristics of politicians on various outcomes (Bhalotra and Clots-Figueras, 2014; Clots-Figueras, 2012).

We chose to analyze outcomes in education because of its importance to development, availability of disaggregated data and the ability of state level leaders to impact education outcomes. Ensuring quality education for all is a Sustainable Development Goal (SDG) and education is a component of the Human Development Index. Education is one of the basic pillars to equip children to avail of various opportunities in life, build a productive labor force and could be improved substantially in India. Given this we expect competent leaders to focus their attention on improving basic education outcomes. Also comprehensive and reliable data disaggregated at district and lower levels are available only for education among the major development outcomes. State-level political leaders have the power to impact education outcomes through lobbying for funds for schools in their districts[7], implementing education schemes, monitoring education policy implementation and influencing bureaucracy. Clots-Figueras (2012) has shown that female leaders at state level are able to improve education outcomes more than male leaders using similar identification techniques. Bhavnani and Jensenius (2015) also study impact of state level leaders from ruling party on literacy over the long term. India is also an important case study to explore this issue due to the recent education mandate that disqualifies citizens with low formal education levels from contesting local body elections and the debate around the issue. Besides, enrollment has seen substantial gains over the last decade across India, but quality of education as measured by learning outcomes is low, declining and varies substantially across the states. Hence we include not only school participation, but also quality of education as outcome in our analysis.

We find no evidence that educated leaders are more competent and deliver better education outcomes for their constituents. Having a college graduate leader does not improve learning outcomes for children, school enrollment or attendance rates. College graduate politicians do not lead to better school infrastructure or more school grants. The null effect is revealed through statistical insignificance of the point estimate and the effect size, adjusted for the precision of the point estimate, remains negligible across OLS and 2SLS specifications. The results are also robust to a

---

[7]An average state in India spends about 16 percent of its total budgetary expenditure on education, with substantial variation across states from 10 to 23 percent, as revealed by recent statistics released by the Reserve Bank of India.

number of specifications, alternative definition of key variables and different definitions of close elections. We also show that the results are not specific to the leader's education cutoff. In the baseline model we compare college graduate and non-graduate leaders, but in robustness analysis we show that even when we compare leaders who have completed secondary or senior-secondary education with those who have not, the results do not change. Since education outcomes might take time to change we also examine impact of leaders from multiple years (up to past four years), but find no significant impact of educated leaders. We also test for heterogeneity of impact across different sub-samples. We test if districts with different levels of urbanization or adult education achievement or poverty levels have varying impact, but overall educated politicians in most cases do not perform any better than less educated politicians. We also test if the impact varies across individual and household characteristics (gender of the child or mothers education level or economic status of the household), but again we find no consistent impact of educated politicians.

Why do educated politicians not provide better education outcomes for children in their constituencies than less educated politicians? It might be the case that formal education has little to do with ability of leader to empathize and address the concerns of people. Also there are avenues other than formal education to develop leadership qualities - through experience and grass-roots work. Understanding constituents' issues and lobbying on behalf of them can be done successfully even without formal education.

These results should of course be interpreted with caution given a range of caveats with any such study. We only examine impact of education of leaders on education outcomes, and it maybe the case that educated politicians use their competence in other spheres like technical education, attracting investments or improving infrastructure to deliver better outcomes. We only analyze impact of state leaders but education of the leader might have a different impact at local or national level. This would not be captured in our analysis. It might be the case that it is difficult for leaders to impact education outcomes, even though they have role to play in state policy making, directing funds to schools in their districts and have some control over education bureaucracy. Limitations of data on education outcomes (which are identified only at the district level) necessitate that we aggregate the impact of multiple leaders at the district level to measure impact; and cannot measure impact directly at the individual constituency level for which they are elected. Our results might be influenced by this averaging at the district level or the other caveats we discuss in the conclusion. More data and research is required to study if educated leaders perform better in other spheres and across levels.

These results, even with all their limitations, suggest that educated leaders are not necessarily better as assumed by policy makers behind the education mandate or the literature which uses education as a proxy for quality. Our findings have direct implications for policy makers. Our results suggest that restricting choice of voters to only leaders who are formally educated might not lead to higher quality leaders or better development outcomes. More research is required to conclusively determine if formal education leads to better leaders, but without that restricting choice of voters might not be advisable. Our findings could serve as inputs if the Supreme Court of India's constitutional bench [8] reviews its original decision upholding the restrictions on candidates or if lawmakers consider a pending bill in the Parliament to overturn the restrictions. It could also lead to rethinking on part of other state governments and policymakers who are planning to introduce such restrictions in their states. Going beyond that our results advance the literature on quality of leaders by being the first to systematically analyze the link between politician's education and competency in India.

The rest of the paper is organized as follows. Section provides a background by linking the current context with the existing literature and providing a conceptual framework for analysis. Section describes how different data-sets have been collated for the purpose of our analysis. The empirical model including the method used for identification is illustrated in Section . In Section , we report the results of our analysis and also various robustness tests and heterogeneity analysis. Section concludes and summarizes limitations of our results.

5.2 Background

5.2.1 Identity of Political Leader

Educated politicians might have a differential impact on children's education only if the identity of the political leader matters for policy making. In a world where candidates can fully commit to implementation of a specific set of policies when elected and care about getting reelected, then politicians' decisions would only reflect the preferences of the electorate (Downs, 1957). In this setting the characteristics of the individual person who wins elections will not matter. Thus, education level, gender, religion, caste, involvement in criminal activities or other aspects of identity

---

[8]Original decision to mandate minimum education for candidates was given by a two member bench which can be reviewed by a full constitutional bench.

of the politician would not affect development outcomes. But citizen-candidate models (Besley and Coate, 1997; Levitt, 1996; Osborne and Slivinski, 1996) have suggested that complete commitment to policy is not possible and the identity of the politician has an influence on the actual policies that are implemented. In this model voters take into account both policy preferences and other relevant characteristics as competency when casting their votes.

Empirical evidence also shows that identity of the politician matters for policies. Jones and Olken (2009) and Besley et al. (2011) use random leadership transitions at national level to show that individual characteristics of the leader matter for economic growth of the country. Extensive literature on India have shown that gender, religion, caste and other characteristics of the political leader play an important role in determining policy outcomes in the fields of education, health, economic growth and public infrastructure (Asher and Novosad, 2013; Bhalotra and Clots-Figueras, 2014; Bhalotra, Clots-Figueras, and Iyer, 2013; Bhalotra, Clots-Figueras, Cassan, and Iyer, 2014; Burchi, 2013; Chattopadhyay and Duflo, 2004; Clots-Figueras, 2011, 2012; Ghani, Mani, and O'Connell, 2013; Halim, Yount, Cunningham, and Pande, 2016; Iyer, Mani, Mishra, and Topalova, 2012; O'Connell, 2015; Prakash, Rockmore, and Uppal, 2014).

5.2.2 Conceptual Framework

We describe a simple model that formalizes our intuition on the impact of educated leaders and guides our empirical analysis. We highlight two important mechanisms through which an educated leader could have an impact. Both of the mechanisms we suggest are unobservable and untestable directly in our empirical framework. However, we flesh out the implications of the two mechanisms and test the implications. The model considers only the action of an elected leader and do not explicitly take into account the actions of voters or other political actors.

We assume that there is a limited amount of resource in terms of influence, effort, or time available to each leader. A leader can use the resource to act upon n different areas such as public infrastructure, health, education, extracting revenue to increase one's own wealth, raising funds for next campaign etc. The leader has to decide on how to distribute the total amount of resource on different activities with the goal of maximizing own utility. Leaders differ among each other on basis of the level of education (G) – which for instance can be whether college graduate or not. If the leader puts e_i level of resource in the ith area (where $i = \{1, ..., n\}$), then the resulting output is given by the following function:

$$A_i(e_i, G) = \Gamma_i(G)\, e_i$$

which can be viewed as an achievement of the leader in the context of this particular area. The choice of this specific functional form is to keep the illustration simple. $\Gamma_i(G)$ is a parameter that captures the competency of the leader in improving outcome i, and it depends on leader's own education level G. It is plausible to assume that $\frac{\partial A_i}{\partial e_i} = \Gamma_i(G) > 0$. If the outcome of interest is education, then the achievement might be increasing enrollment rate, learning levels of students or improving school quality. We observe the achievements but do not observe the actual resource allocated or effort exerted by the leader across different areas directly.

The leader derives utility from achievement on various items through several channels. Firstly, achievements reflect a leader's performance which in turn affects the probability of winning the next election. Secondly, leaders also get direct utility from achievement on areas that create a lasting legacy or give them a sense of satisfaction and fulfillment. Therefore the leader's utility is a function of his/her achievements on the n different items. For ease of exposition we assume a log-linear form of the utility function, and express the leader's problem as follows:

$$\underset{\{e_1, e_2, ..., e_n\}}{\text{maximize}} \quad \sum_{i=1}^n \Theta_i(G) \log[A_i(e_i, G)]$$
$$\text{subject to} \quad \sum_{i=1}^n e_i = 1$$

The preference parameter $\Theta_i(G)$ reflects how much weightage is given to achievement in the ith area by a leader with G level of education. We assume $\Theta_i(G) \geqslant 0$ and $\sum_{i=1}^n \Theta_i(G) = 1$. The constraint indicates that there is an upper bound to the total resource available and it has been normalized to 1.

Solving the first order conditions of this problem we find that the leader's optimal allocation of resource towards the ith item is given by: $e_i^* = \Theta_i(G)$. Therefore the optimal level of achievement is expressed as:

$$A_i(e_i^*, G) = A_i^*(G) = \Gamma_i(G)\, \Theta_i(G)$$

So the optimal achievement on ith item is directly proportional to competency and preference of the leader towards this particular item. Since both competency and preference parameters are non-negative, therefore achievement will increase with G if at least one of competency ($\Gamma_i(G)$) and preference ($\Theta_i(G)$) increases with G. In other words, an educated leader yields better outcome in an area if he/she has higher competency or higher preference towards that area as compared to a

non-educated leader.

A competent leader or leader who understands a particular issue might be able to achieve better outcome in the area i than a less competent or less knowledgeable leader for the same level of effort. Leaders might also differ in their achievements in sector i because they obtain more utility from achievements in the sector above the ones obtained by other leaders due to their preferences. While we cannot empirically separate out preference from competence, we test the claim in our empirical model that achievements of educated leaders are higher than less-educated leaders.

Besley and Reynal-Querol (2011) argue that good leaders might differ based on competency and prefer broad-based policies and advocate public goods and infrastructure that generate wide benefits. They argue that educated leaders are better citizens, have higher talent and higher concern for social welfare. Elementary education is a basic public good whose widespread provision can lead to improvement in welfare of citizens. If the arguments made by Besley and Reynal-Querol (2011) or Atkinson et al. (2016) are valid then educated politicians would lead to better quality of elementary education as they are more competent and would prefer to invest their resources in education due to its impact on long term welfare of citizens. Even if educated leaders do not prefer to invest in education, their higher competency by itself should on average lead to better education outcomes.

While acknowledging these arguments Carnes and Lupu (2016) argue that the link between formal education, competence and leadership is not as straightforward. It is possible to gain human capital and skills required for being an effective leader without formal education. Formal education does not just reflect human capital but also the privileges of being able to obtain education. Societal restrictions and economic hardship might hinder a talented person from obtaining formal education while a less talented person might be able to obtain formal education. There still exists widespread discrimination, though declining, in obtaining even elementary education by women and individuals from lower castes in India.

Carnes and Lupu (2016) also argue that human capital obtained through formal education alone does not necessarily improve quality of leaders. Crucial factors like character, personality, ability to listen and understand people's grievances and other leadership qualities are not the focus of formal education and could be obtained without it. Even basic reading and arithmetic skills could be obtained outside of formal schooling. In instances where specialized skills are required leaders could rely on qualified bureaucracy to help formulate solutions. Several studies have also shown that most qualified and those with highest grades are not necessarily most successful people in the

society (Gottesman and Morey, 2006; Clotfelter et al., 2007). Several studies in the 1960s and 70s and recent followups by Carnes (2012, 2013) which studied relationship between education level of leaders and their attitudes, choices and decisions when in office found no difference in behavior between more and less educated politicians (see Carnes and Lupu (2016) for details).

The empirical evidence on impact of education level of the leader on policies is thin and mixed. Dreher et al. (2009) find that professional and highly educated leaders are more likely to implement market-liberalizing reforms. On the other hand, Carnes and Lupu (2016) show that across contexts and wide range of outcome indicators politicians with a college degree perform the same or worse than non-college graduate politicians. Educated leaders at national level in a cross-country database, legislative leaders in the US and local municipal leaders in Brazil all perform no better than non-educated leaders. This holds across a range of outcomes including economic growth, inequality, social unrest, interstate conflict, unemployment, inflation, reelection, legislative productivity, and corruption.

Most of the studies which found a positive impact of education qualification of the leader analyze leaders at the national level but the impact of education of the leader might be different at lower administrative levels (except a few studies such as Martinez-Bravo (2014)). Leaders at district or constituency level work more as "fixers" or lobbyists for their constituents and have smaller role to play in broader policy making (Chopra, 1996). Analyzing the activities of elected state constituency representatives across Indian states, Jensenius and Suryanarayan (2015) find that state leaders spend most time in their constituencies rather than debating legislation in state assemblies. The time spent by leaders debating legislation and making policies has also been declining over time – from 45 days per year in 1967 to about 34 days per year – and most legislation is passed without much debate. Politicians tend to spend most of their time in their constituencies addressing their constituents' complaints, attending social functions, being part of local government bodies, helping individuals in accessing various government schemes, lobbying the district and state administration to implement their favored schemes and use their networks to attract investment to their constituency (Chopra, 1996; Jensenius, 2015). Formal education might be helpful but is not essential to gain the skills required for this set of tasks.

5.2.3 State leaders and education policy

India has a federal structure with significant power devolved to the states and local level. For elections, states are divided into single- member constituencies in which candidates are elected in first-past-the-post elections. Constituencies within each state are drawn up so to ensure the same number of people in each of them, though they vary in size across states. District is the administrative unit below the state and consists of multiple constituencies, with a district having an average of nine constituencies.

Education is on the concurrent list in the Indian constitution – where both the federal and state governments have jurisdiction, but state governments play a major role in education policy at the primary and secondary level. Legislators can influence state policy by participating in debates, and influencing other legislatures at the state level. They can also direct funds to the their districts' educational office and influence policy implementation through their participation in local government bodies (Singh and Cruz, 1997). Legislatures can monitor school infrastructure and progress, lobby the state government for funds to open new schools or help in accessing existing grants or programs for schools from state or federal government. State legislatures also have control over state bureaucracy through role in promotions and job assignments/transfers (Krishnan and Somanathan, 2013; Nath, 2015; Sukhtankar and Vaishnav, 2015; Asher and Novosad, 2013). This influence can be used by legislatures to push their specific policy priorities and also demand results. State legislative representatives can use their discretionary development funds for any development work. Empirical evidence also links various characteristics of state legislatures with differences in policy outcomes in various spheres (Bhalotra and Clots-Figueras, 2014; Bhalotra et al., 2013, 2014). Clots-Figueras (2012) shows that districts with higher proportion of women state legislative leaders have higher primary school completion rate than other districts. She argues that women leaders prefer to invest in children's education and are able to influence district policy outcomes. If educated leaders are more competent then they can, if desired, presumably improve education outcomes for their constituents.

In this paper we focus our analysis on determining whether educated political leaders have an impact on the schooling infrastructure and allocation of funds to schools. In addition, we consider children's enrollment and learning as other primary indicators of education outcomes. With rapid expansion of elementary education in the last decade enrollment rates in primary schools have reached saturation levels (>95 percent enrollment across India), but learning outcomes among children have stayed low and declined in recent period. According to Annual Status of Education

Report (ASER) only 40 percent of children in class III can read a class I level text and only 26 percent of children in class V can do subtraction in 2014. These levels have declined from 49 and 43 percent respectively in 2007 (ASER 2014). Basic knowledge and skills - not enrollment and years of schooling - are key to empower children to realize their potential in life (Pritchett, 2013). This has been recognized widely by non-government organizations (NGO) and state governments across the country [9]. Pratham, a large education NGO, which first highlighted these issues has implemented several initiatives to improve learning outcomes since 2007. Pratham working in collaboration with state and local governments has organized short-duration intensive learning camps to improve basic skills for over 430,000 children in 2014-15. Randomized evaluations of learning outcome improvement programs implemented by Pratham in collaboration with local governments in states of Bihar, Uttarakhand, Haryana, Gujarat and Maharashtra have found substantial improvements in reading and math skills among children (Banerjee, Cole, Duflo, and Linden, 2005; Banerji, Berry, and Shotland, 2013; Banerjee, Banerji, Duflo, Glennerster, and Khemani, 2010). These programs are simple, low cost interventions and usually involve village volunteers or hired young adult helping teachers 1-2 hours a day or through intensive learning camps to focus on teaching core competencies which are supposed to be taught in first and second grade. Several state governments have also implemented their own programs to improve reading and arithmetic skills. If political leaders wish to improve learning outcomes or are more competent than other leaders, improving learning outcomes is not a far-fetched goal with the resources they have at their disposal.

5.3 Data and Summary Statistics

The empirical analysis investigates the causal relationship between being represented by college graduate politicians versus a non-college graduate politicians and different education outcomes (learning and enrollment outcomes for individuals of school age and school quality outcomes) when those politicians were in power in their districts.

---

[9]The central government in its 2017 budget acknowledging the importance of learning outcomes has decided to make quality assessment a central plank in its elementary education policy.

Education Data

We create a dataset that combines information on politicians standing in state assembly election in India with different education outcomes. Education outcomes are from different data sources and both at individual child and school level. The individual child outcomes are learning outcomes for reading, mathematics and english from Annual Status of Education Report (ASER) and enrollment outcomes from National Sample Survey (NSS). School quality outcomes are obtained from annual census of schools - District Information System for Education (DISE). We use multiple data sources as the range of outcomes are not present in any single dataset and it also checks robustness of our results to source of data.

ASER Data

ASER, an annual district representative survey, documents children's schooling status and basic learning levels in all rural districts in India. The survey has been conducted every year from 2005 to 2014 by a group of over 30,000 trained volunteers from over 700 partner organizations under the leadership of Pratham, an educational NGO. The survey is conducted between September-November, and covers a random sample of 20 households in 30 villages in each of India's rural districts (550) totaling about 300,000 households across the country each year and approximately 600,000 children in the age group of 3-16.

ASER tests all children in the household between the ages of 5 and 16 for basic arithmetic and basic reading proficiency in the vernacular language using rigorously developed testing tools[10]. The same test is given to all children across the years. The reading assessment has four levels: letters, words, a short paragraph (a class 1 level text), and a short story (a class 2 level text). Similarly, the arithmetic assessment consists of four levels: single-digit number recognition, double-digit number recognition, two-digit subtraction with carry over, and three digit by one digit division (corresponding to what students are expected to know in grade 3 or 4). These levels are converted into a continuous scale of 0-4 in our main analysis. The highest level for which children are comfortable is marked. In years 2007, 2009, 2012 and 2014 children were also tested for their competency in basic English. In addition to learning outcomes, basic household information (household size, parental education and some information on household assets) and village infrastructure information (existence of electricity, permanent road, ration shop, bank, schools and

---

[10]The tools are available at http://www.asercentre.org/p/141.html

143

health facilities) is also collected. In 2007 and every year since 2009 ASER has also collected data on school infrastructure, enrollment, attendance and fund flows from one government primary school in each surveyed village.

We use data for the years 2006 to 2014 for children between 6 to 16 years of age. We divide our sample into two groups – children between 6 to 10 years of age which corresponds to primary school age (grade 1 to 5) in India and children between 11 to 16 of age corresponding to middle and secondary school age (grade 6 and above) to study impact of educated politicians on their learning outcomes. Overall our sample includes 1.29 million children between 6-10 years of age and 1.3 million children between 11-16 years of age. Though enrollment in India has increased over time learning levels have declined.

NSS Education Data

We use four rounds of National Sample Survey (NSS) data to obtain data on enrollment of children in both urban and rural areas. The data comes from Education rounds of NSS conducted in 2007-08 and 2014 and Employment and Unemployment surveys from 2009-10 and 2011-12. All four surveys record information on enrollment status in education institutions for all members of the household and we use the same age groups as used in ASER data analysis [11]. Overall our sample from NSS includes 145,000 children aged 6-10 years of age and 170,000 children aged 11-16 years of age. We use NSS surveys for enrollment outcomes as it covers both rural and urban areas, as opposed to only rural areas in ASER. But NSS surveys do not collect information on learning outcomes. The outcome variable used in the analysis is a dummy indicating if child is not enrolled in school. Overall 6.9 percent of 6-10 year old children and 13.5 percent of 11-16 year old children are not enrolled in school in the NSS sample (Table 5.1).

DISE Data

DISE is a database on all schools in India with information on school infrastructure, funding, enrollment, teachers and other aspects of schools maintained by the Government of India. This is updated annually and cross-validated. DISE has data on over 1 million schools across India

---

[11]Though the sample size, sampling stratergy and options present in the enrollment question vary to some extent across surveys, there is enough commonality in the basic question that they could be compared. The pattern in enrollment rate (of increase across years) is consistent with other secondary data.

Table 5.1: Summary Statistics from Child Level Data

Variables	6-10 age-group			11-16 age-group		
	Obs.	Mean	SD	Obs.	Mean	SD
	(1)	(2)	(3)	(4)	(5)	(6)
ASER Data						
Reading score	1,294,549	2.071	1.388	1,303,674	3.425	1.047
Math score	1,279,657	1.893	1.213	1,295,643	3.123	1.076
English score	567,828	2.678	1.386	576,782	3.959	1.270
Proportion not enrolled in school	1,266,493	0.016	0.124	1,302,898	0.060	0.238
Age of child	1,294,549	8.113	1.428	1,303,674	13.291	1.624
Proportion of female children	1,294,549	0.467	0.499	1,303,674	0.478	0.500
Proportion whose mother went to school	1,250,366	0.534	0.499	1,261,765	0.499	0.500
Household size	1,294,549	6.553	2.858	1,303,674	6.369	2.714
Proportion living in katcha houses	1,214,531	0.368	0.482	1,228,590	0.338	0.473
Proportion living in semi-pucca houses	1,214,531	0.305	0.460	1,228,590	0.306	0.461
Proportion living in pucca houses	1,214,531	0.328	0.469	1,228,590	0.356	0.479
Proportion in households with television	1,203,449	0.466	0.499	1,218,052	0.506	0.500
Proportion in households with mobile phone	1,197,277	0.650	0.477	1,212,061	0.695	0.461
Household's with electricity	1,022,433	0.645	0.478	1,054,933	0.679	0.467
Proportion in villages with electricity	1,207,515	0.909	0.287	1,221,079	0.927	0.261
Proportion in villages with pucca road	1,202,562	0.746	0.435	1,216,265	0.766	0.423
Proportion in villages with ration shop	1,200,705	0.703	0.457	1,214,089	0.718	0.450
Proportion in villages with a bank	1,199,046	0.240	0.427	1,212,105	0.258	0.438
NSS 2004-05 Data (based on district level aggregates)						
Proportion of female population	1,258,269	0.487	0.027	1,271,747	0.488	0.028
Primary completion rate - adult male	1,258,269	0.651	0.150	1,271,747	0.659	0.147
Primary completion rate - adult female	1,258,269	0.374	0.199	1,271,747	0.384	0.198
Proportion of ST	1,258,269	0.156	0.266	1,271,747	0.148	0.254
Proportion of SC	1,258,269	0.197	0.122	1,271,747	0.201	0.122
Proportion of OBC	1,258,269	0.410	0.237	1,271,747	0.411	0.234
Proportion of other caste	1,258,269	0.237	0.204	1,271,747	0.241	0.206
Proportion of urban population	1,260,105	0.179	0.150	1,273,831	0.187	0.151
NSS Education Data (Various rounds 2007 to 2014)						
Proportion not enrolled in school	145,297	0.069	0.253	172,682	0.135	0.342
Proportion of female children	146,363	0.464	0.498	174,457	0.464	0.498
Proportion living in urban areas	146,363	0.349	0.476	174,457	0.371	0.483
Household size	146,363	6.145	2.618	174,457	5.893	2.447
Proportion Hindu	146,350	0.732	0.442	174,435	0.733	0.426
Proportion Muslims	146,350	0.164	0.370	174,435	0.153	0.360
Proportion ST	146,319	0.134	0.340	174,397	0.132	0.338
Proportion SC	146,319	0.178	0.382	174,397	0.174	0.379
Proportion OBC	146,319	0.415	0.492	174,397	0.406	0.491

Source: Author's calculation using ASER and NSS data. The unit of observation is children of different age groups.

Table 5.2: Summary Statistics of District Level Variables from DISE Data

	Mean	SD	Obs
Free Textbooks	0.82	0.15	2611
Furniture for Students	0.67	0.24	2611
Atleast one Girl's Toilet	0.82	0.21	2611
Electricity Connection	0.53	0.35	2611
Safe Drinking Water	0.76	0.22	2611
Library	0.56	0.28	2611
Playground	0.58	0.22	2611
Computer	0.27	0.23	2611
PCA school facilities	0.02	1.85	2611
Annual average school grants	18793	14092	2695

Notes: The unit of observation is district. The sample corresponds to the full sample used in DISE regressions.
Source: Authors' calculation from DISE data.

for each year; and is the most comprehensive source of data on schools in India. We use data for five years from 2010 to 2014 on school infrastructure and funding outcomes. We aggregate this data on various outcomes of interest at the district level. Table 5.2 presents the summary statistics from DISE data. We use Principle component analysis (PCA) to generate an index of school infrastructure and use it as one of the outcome variables. The index is the first component of the PCA generated using data on free textbooks, furniture, girl's toilet, electricity connection, access to safe drinking water, library, playground and computers. The second outcome variable is the average grants received by schools in the district in one year. There is substantial variation in both outcome various across districts.

Political Data

India is a federal republic with parliamentary system of government at the state and the national level. Several powers are devolved to the state, district and village level government. Each state has a legislative assembly (state governments) which plays a big role in educational policies and expenditures, especially at the primary and secondary level. States are divided into districts which are important administrative units for various decisions and districts in turn are divided into single-member constituencies in which candidates are elected in first-past-the-post elections. Each assembly constituency is designed so as to have almost the same number of inhabitants within the state. On average each district has about nine constituencies. The term of each elected state representative is five years, unless the assembly is dissolved before end of its term.

Table 5.3: Education Qualification of Candidates in Elections

Education of the candidate	Sex of the candidate								
	Female			Male			Total		
	Num	Col %	Cum %	Num	Col %	Cum %	Num	Col %	Cum %
Illiterate	7	0.5	0.5	35	0.2	0.2	42	0.3	0.3
Literate	49	3.7	4.2	202	1.4	1.6	251	1.6	1.8
5th Pass	37	2.8	7.0	292	2.0	3.6	329	2.0	3.9
8th Pass	80	6.0	12.9	726	4.9	8.5	806	5.0	8.9
10th Pass	191	14.3	27.2	2147	14.6	23.1	2338	14.5	23.4
12th Pass	184	13.8	41.0	2341	15.9	39.0	2525	15.7	39.1
Graduate	291	21.8	62.8	4011	27.2	66.2	4302	26.8	65.9
Graduate Professional	131	9.8	72.6	2252	15.3	81.4	2383	14.8	80.7
Post Graduate	312	23.3	95.9	2414	16.4	97.8	2726	17.0	97.6
Doctorate	55	4.1	100.0	323	2.2	100.0	378	2.4	100.0
Total	1337	100.0		14743	100.0		16080	100.0	

Notes: Table based on information about winner and runner-up candidates. A person who has a Bachelor's
degree or diploma is considered a Graduate in India. Source: Author's calculation using ECI-ADR data.

A dataset on politicians contesting state assembly elections in India between 2004 and 2014 was constructed using information obtained from the Election Commission of India (ECI) and the Association for Democratic Reforms (ADR). The ECI provides data on the number of votes, gender and party affiliation of all winner and runner-up candidates for all state assembly elections in India. Following a 2003 Supreme court judgment all individuals contesting elections have to file an affidavit with the election commission listing their education level, assets, criminal cases among other details. The Association for Democratic Reforms (ADR) has scanned all these affidavits and provided the information online for central and state elections from 2004 onwards. We combined this information from ECI and ADR to construct a detailed portrait of all winner and runner-up candidates for state assembly elections from 2004 to 2014 . Among winner and runner-up candidates about 39 percent have not completed college education (Table 5.3). Table 5.4 presents summary statistics from the electoral data. About 38 percent of districts had at least one close election between graduate and a non-graduate and almost equal number of close elections are won by both types of candidates (49 and 51 percent respectively).

Merged Data

ASER and NSS rounds provide information on residence of a child only at the district level whereas political leaders are elected at the constituency level, which are below the district level. To merge the education and political datasets we aggregated the election data at the district level. For each child in the sample politicians who were in power in the year when the child's enrollment and

Table 5.4: Summary Statistics of District Level Variables from Election Data

Variables	Obs.	Mean	SD
	(1)	(2)	(3)
Proportion of seats won by graduates	1,127	0.594	0.251
District with at least one graduate leader	1,127	0.953	0.212
Proportion of seats won by graduates in close elections against non-graduates	1,127	0.0394	0.0979
District with at least one graduate leader who won in close election against non	1,127	0.201	0.401
Proportion of seats with close election between graduate and non-graduate	1,127	0.0814	0.138
District with at least one close election between graduate and non-graduate	1,127	0.378	0.485
Proportion of seats with election between graduate and non-graduate	1,127	0.407	0.249
District with at least one election between graduate and non-graduate	1,127	0.876	0.330

Notes: The unit of observation is district in an electoral year. The sample corresponds to the full sample used in child level regressions. Close election is the one where the winner beat the runner up by less than 3 percent of votes.

Source: Authors' calculation from ADR and ECI data combined.

learning levels were recorded and the two years prior to it are identified. ADR has data on candidate's education levels only for state elections conducted after 2004, when the law requiring candidates to publically report their education level came into force. Since the elections for state assemblies in different states are conducted in different years, the starting year of our data differs for each state - the first election year after 2004 for that state. For example, for Maharashtra which had election in 2004 we have data on politician's education level for all years from 2004 to 2014, but for Gujarat our data starts only in 2008 the first election year in Gujarat after 2004. Since we are using average value of three years for all political variables in our base specification the data used in the estimation for a state start two years after the start of election data for the state. For example, for all districts in Maharashtra our base estimation uses ASER data starting from 2006 whereas for districts in Gujarat estimation data starts only in 2010. A child living in rural areas of Jalna district of Maharashtra surveyed in 2007 will be impacted by politicians elected in the district from 2005 to 2007.

5.4 Empirical Strategy

In this section we lay out an econometric model to identify the causal effect of having a college educated versus non-college educated leader on the educational outcomes of children. As explained in the previous section, the data on leaders' education can be merged with the data on outcome variables only at the district level. Therefore our treatment variable is defined for each district, as the fraction of constituencies with an educated leader in the district. When the outcome is mea-

sured at the child level using ASER or NSS data (learning outcomes, enrollment etc.), then our unit of analysis is a child because we also want to control for child and household specific covariates. On the other hand, when we use DISE data on schools, the outcomes are measured at the district level (average school quality, average grants etc.) and hence district itself is the unit of analysis. However, in both these cases, the treatment variable is defined at the district level for a given year. The following equation illustrates our basic specification which is estimated by the Ordinary Least Squares (OLS) method:

$$Y_{idst} = \alpha_{ds} + \sigma_t + \beta G_{dst} + \gamma \mathbf{X}_{idst} + \delta_t \mathbf{Z}_{ds} + \eta_{st} + \varepsilon_{idst} \tag{1}$$

The education outcome of child i living in district d of state s in year t is given by Y_{idst}; and G_{dst} is the fraction of assembly constituency seats in the district held by a college-educated politician during the last three years.[12], [13] District specific time invariant unobserved heterogeneity is taken into account by including district fixed effects α_{ds}. Since districts are nested within states, therefore the district fixed effects also subsume the state fixed effects. The year fixed effects (σ_t) consider the overall changes in the economy including the impact of growth and various nationwide educational policies. Several observable characteristics at the level of child, household and village are included in the vector $\mathbf{X}_{idst}$. Child level covariates are gender and age-cohort specific dummy variables, and an indicator of whether child's mother attended school. Household level variables include household size, square of household size, dummy variables indicating the structure of the house[14], ownership of assets (television and mobile phones), and whether use of electricity was observed on the date of survey. The village specific variables capture access to electricity, paved (*pucca*) road, ration shop and bank. While the district fixed effects control for regional characteristics that do not change over time, there are factors such as demand for education or level of development that vary over time. Some of these district specific time varying effects are taken into account by interacting the year fixed effects with measures of baseline characteristics given by the vector $\mathbf{Z}_{ds}$. Using National Sample Survey data of 2004-05, district level sex ratio, adult primary education completion rates for males and females, caste composition in the population, and proportion of urban population are included in $\mathbf{Z}_{ds}$. Most of the education policies are imple-

---

[12]Following Clots-Figueras (2012), the main specification considers the average fraction of seats held by a college educated politician over the past two years and the current year. In the robustness section, we consider alternative lag periods to calculate this variable, and our results remain unchanged in those specifications.

[13]A college educated individual with at least a Bachelor's degree or diploma is referred to as graduate in the India.

[14]Depending on the building material, the structure of the house is measured by the following three categories, arranged in the increasing order of housing quality: *katcha*, *semi-pucca* and *pucca*. *Katcha* is considered as the base category in the regression.

mented by the respective state governments; besides, there is a great deal of heterogeneity across Indian state economies which are likely to follow very different trajectories of development. We control for all time varying state level factors by including state specific year fixed effects (η_{st}) in the regression.

The main challenge of identifying β from Equation 1 is the possibility that some omitted variable which varies across districts and over time, may be correlated with both G_{dst} and Y_{idst}. While the district fixed effects take care of inherent differences that do not change over time, the presence of time varying unobservable effects at the district level cannot be ruled out. For instance, in regions that have experienced higher growth of educated individuals, voters' preference for education may be manifested through higher propensity to elect college educated politicians in recent elections than in past elections. The fraction of seats held by college educated leaders may be endogenously determined due to the presence of such unobservable factors.

Identification

To tackle the endogeneity problem, we use the fraction of seats won by college educated politicians in *close* elections between a college educated and a non-college educated politicians (GC_{dst}) as an instrument for the overall fraction of seats held by college educated leaders (G_{dst}). Close elections are defined as those where the margin of victory is small. For the main specification, we consider an election to be close when the winner beats the runner-up by less than 3 percent of total votes, and measure the instrument accordingly.[15] Insofar as the vote difference between the top two candidates in an election is arbitrarily small, the winner will be determined by chance; hence the use of close election provides a plausible basis for constructing the instrument in this context. We provide further discussion on the validity of instrument in a subsequent section.

This empirical strategy has been used in the literature to identify impact of other personal characteristics of leaders (e.g. gender, religion) on various development outcomes (Bhalotra and Clots-Figueras, 2014; Bhalotra, Clots-Figueras, Cassan, and Iyer, 2014; Clots-Figueras, 2011, 2012). Identification in this method relies on the quasi-randomness of the outcome of a close election. The Indian electoral system follows the first-past-the-post voting system where the candidate who

---

[15]The margin of victory is defined as the difference in the share of votes between the winner and the runner-up, where the total turnout is used as the denominator to calculate the vote shares. In the robustness section, we use various other levels of margin of victory to define close elections between a college educated and a non-college educated candidate.

gets more votes than any other candidate wins the election. The probability that a candidate will win is a function of the margin of votes between the winner and the runner-up, and this probability changes discontinuously at the point where the margin of votes is zero. Considering those elections where the contest takes place between a college educated and a non-college educated politician, in an arbitrarily small neighborhood around this point of discontinuity, the constituencies which elect a college educated leader versus those which elect a non-college educated leader are similar in all characteristics except the education level of the leader.[16] Hence this discontinuity at zero margin of votes is essentially similar to random assignment of treatment. Since the main explanatory variable is at the district level, we aggregate over the constituency specific discontinuities in treatment assignment within district; thus we have a fuzzy regression discontinuity design in our empirical set up. The model is estimated through a two-stage least squares (2SLS) method and is given below:

$$Y_{idst} = \alpha_{ds} + \sigma_t + \beta G_{dst} + \lambda TC_{dst} + \sum_{j=1}^{N} \pi_j I_{jdst} \times F(M_{jdst}) + \sum_{j=1}^{N} \mu_j I_{jdst} + \mathbf{X}_{idst}\gamma + \mathbf{Z}_{ds}\delta_t + \eta_{st} + \epsilon_{idst}$$
(2)

$$G_{dst} = \omega_{ds} + \nu_t + \theta GC_{dst} + \rho TC_{dst} + \sum_{j=1}^{N} \phi_j I_{jdst} \times F(M_{jdst}) + \sum_{j=1}^{N} \psi_j I_{jdst} + \mathbf{X}_{idst}\xi + \mathbf{Z}_{ds}\zeta_t + \tau_{st} + e_{idst}$$
(3)

Equation 2 is the second stage and Equation 3 is the first stage. The main explanatory variable G_{dst} which is potentially endogenous, is instrumented by the proportion of educated leaders who win in close elections against a non-college educated candidate GC_{dst}. Note that unlike the outcome of a close election, the existence of close election may not be random: it may depend on the number of educated candidates or the prevailing political competitiveness in the district. Therefore we control for the fraction of seats that had close elections between college educated and non-college educated candidates in the district (TC_{dst}). This also captures any direct effect of having close elections, such as greater effectiveness of leaders due to higher political competitiveness in the region. The specification also controls for a third order polynomial in the victory margins of every college

---

[16]In the empirical analysis educated leader refers to a leader who has completed college and obtained a degree or a diploma.

educated versus non-college educated election (close or non-close) in the district. The margin of victory between a college educated and a non-college educated candidate in election j is M_{jdst}. The polynomials, denoted by $F(M_{jdst})$, are interacted with I_{jdst} which indicates the existence of an election between a college educated versus non-college educated politician j in the district during the period considered. We also test if the results are robust to varying degrees of the polynomial function. Our model is based closely on model used to study the impact of women leaders on health and education outcomes by Clots-Figueras (2012) and Bhalotra and Clots-Figueras (2014). The rest of the variables included in the 2SLS analysis are same as in the OLS regression. In both these models, the standard errors are clustered at the district level to allow for any possible correlation in the error terms among observations within the same district.

5.5 Results

In this section, we present results of impact of educated politicians on our main outcome variables of learning levels, enrollment and schools. The dependent variables for learning levels measure cognitive outcomes in terms of standardized reading and mathematics score[17]. On enrollment, we investigate the impact of leaders education on probability of being out of school. The unit of analysis for learning and enrollment outcomes is the individual child. For schooling outcomes, we analyze impact on index of school infrastructure and grants received by the school. All dependent variables for school outcomes from DISE data are averaged at the district level and unit of analysis is the district-year.

For each of the child-level dependent variables, we present three different specifications where control variables are gradually added to investigate if the results remain comparable across these models. The main variable of interest is the fraction of seats in the district held by educated leaders. The common covariates included in all specifications are district fixed effects and year fixed effects. Beginning with the most parsimonious specification that excludes any other covariate, control variables at the level of individual, household, village and district are cumulatively added. The final model also includes state specific year fixed effects to consider time varying unobservables at the state level. For school outcomes only two models for each dependent variable are presented as

---

[17]We use age-wise standardized test scores as a measure of cognitive outcome. For any given age between 6–16, we consider children of that age from all the survey years, and calculate the mean and standard deviation of their test scores in a specific subject, and thus calculate the z-score for that subject. The implication of using these standardized test scores (z-scores) as outcomes is that the magnitude of effect from the regression can be interpreted in terms of standard deviation in test scores.

there are no individual or village level data to use as controls in DISE.

We begin with the results from the OLS specification before discussing the 2SLS results that take into account the potential problem of endogeneity of the main explanatory variable. Table A1, A2 and A3 present the OLS results for reading score, mathematics score and enrollment respectively. Considering that children at the primary and post-primary levels may have very different learning trajectories, separate regressions are estimated for children in the age-group of 6–10 years and 11–16 years.

The OLS regressions show no significant effect of the proportion of college-educated leaders in the district on children's reading and mathematics scores or on probability of being out of school. Only in one of the specifications that does not control for any observable covariates, we find positive effect on reading score of 11–16 year old children. It shows that a 10 percentage point increase in the fraction of college graduate leaders in the district leads to 0.006 standard deviation increase in the reading score of children in this age-group. This effect is significant at 10 percent level. However, this regression does not control for any observable child, household, village, or district level characteristics, and once these control variables are included, the effect becomes insignificant and also reduces in magnitude. Thus, the OLS results do not show any systematic effect of college graduate leaders on enrollment and learning outcomes of children.

Table A5 presents results from OLS regressions for school outcomes. We find a negative significant impact of educated leaders on average grants received by schools in the district. On average, having one additional graduate leader in the district (11 percentage points increase in number of educated leaders in an average district with 9 constituencies) leads to decrease in average grants for schools by Rs.468. The impact is not significant on school infrastructure.

As discussed earlier, these regressions could suffer from endogeneity problem and not capture the causal effect of college graduate leaders on the outcomes, hence in the subsequent analysis we focus on the instrumental variable estimation in a fuzzy regression discontinuity setting.

Validity of Instrument

Before analyzing the effect of graduate leaders on education outcomes, it is imperative to test whether the instrument is a good predictor of the endogenous variable in the first stage of the 2SLS estimation. Corresponding to the reading score regressions, Table 5.5 presents the first stage re-

Figure 5.1: First Stage Illustration: Sample of all Districts with Close Elections
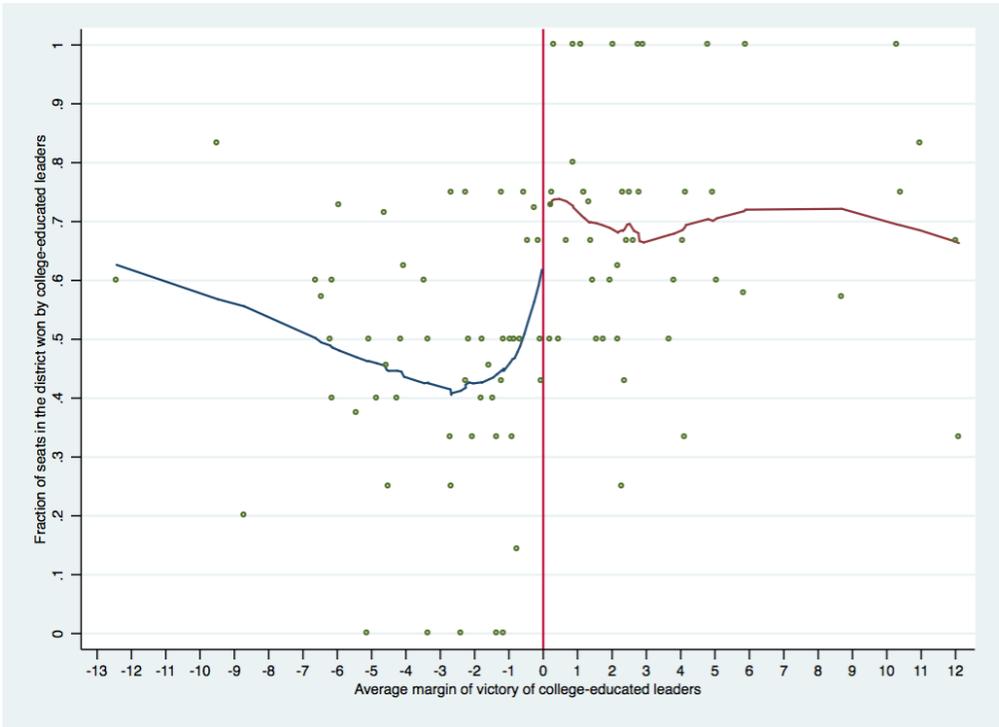
Table 5.5: First Stage of the 2SLS Estimates of the Effect on Children's Reading Score

	Fraction of seats won by a graduate					
	6–10 age-group			11–16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won by a graduate	1.052***	1.035***	0.987***	1.039***	1.018***	0.973***
in close elections	(0.144)	(0.133)	(0.133)	(0.133)	(0.125)	(0.125)
Vote margins: third order polynomial	Yes	Yes	Yes	Yes	Yes	Yes
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Age cohort fixed effects		Yes	Yes		Yes	Yes
Child, household & village controls		Yes	Yes		Yes	Yes
District controls			Yes			Yes
State by year fixed effects			Yes			Yes
Observations	1,262,927	896,820	869,544	1,272,189	928,499	904,149
Number of districts	563	563	545	563	563	545
R-squared	0.387	0.399	0.433	0.392	0.401	0.438

Notes: Robust standard errors clustered at the district level are in parentheses. Close elections are defined as between a graduate and a non-graduate in which the difference in vote share between the winner and the runner up is less than 3 percent. Individual level controls are dummy variables for children's age cohort and gender, and whether mother went to school. Household controls are household size, square of household size, type of building, whether household owns television or mobile phone, and whether use of electricity was observed in the household. Village controls include indicators of whether village has access to pucca road, electricity, ration shop, and bank. District controls are baseline characteristics interacted with year dummies; they include urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas estimated from 2004-05 National Sample Survey data. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

gressions for 6–10 and 11–16 age-groups. The coefficient of the instrumental variable, i.e. fraction of seats won by a graduate leader in close elections, is found to be statistically significant at 1 percent level in all specifications and for both the age-groups. The first stage F-statistics for the instrumental variable ranges between 53 to 60 for 6–10 age-group and 60 to 66 for 11–16 age-group. The point estimate is also stable. Result from the final model for the sample of 6–10 age-group shows that holding the fraction of constituencies with close election constant, a 10 percentage point increase in the fraction of constituencies where a graduate leader won against a non-graduate leader in a close election leads to 9.87 percentage point increase in the overall fraction of seats held by college graduate leaders in the district. The equivalent estimate for 11–16 age-group is 9.73. The results are similar for first stage analysis using NSS education and DISE data (Tables A4 and A6).

We also provide graphical illustration of the first stage. We plot the overall fraction of college-educated legislators against the average vote margin between college-educated and non-college educated candidates across districts. Figure 5.1 uses all the elections in districts with at least one close election (sample on which our identification is based), while Figure 5.2 restricts sample

Figure 5.2: First Stage Illustration: Sample of all Districts with at least one Election between Educated and Less-Educated Candidates
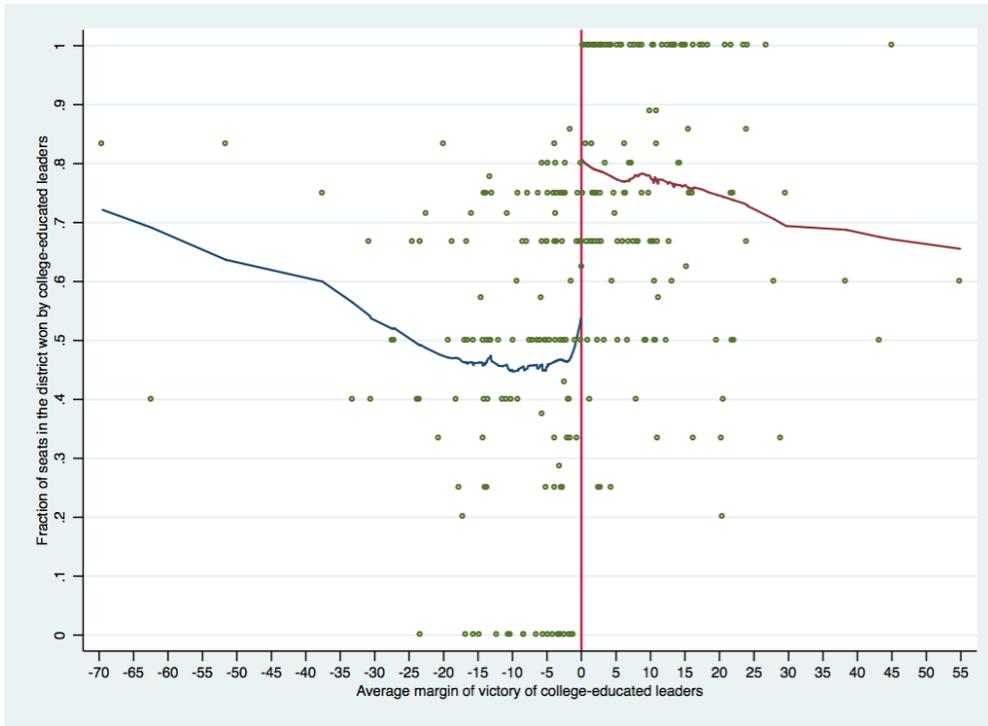
to districts which have exactly one election between college-educated and non-college graduate politicians. We use bin size of one percentage point as suggested by Imbens and Lemieux (2008), and we plot a lowess smoothing line on each side of the discontinuity. A college-educated politican winning a close election in a district increases the fraction of constituencies in the district won by a college-educated politicians by about 10 percentage points (Figure 5.1).

Validity of Close Elections Assumption

We conduct several checks to ensure the validity of our assumption that results in close elections are quasi-random. In summary statistics, we saw that the probability of winning a close election is the same for graduate and non-graduate candidates. But some recent studies have questioned the validity of the close election premise by showing that

Table 5.6: Comparing Candidate and Constituency Characteristics across Close Elections with Educated and Non-Educated Winners

	Non-College educated winners	College educated winners	Difference
Proportion of winners who are women	0.067	0.050	0.017 (0.019)
Proportion of winners with criminal cases	0.33	0.31	0.026 (0.038)
Average number of college-educated candidates	3.55	3.43	0.12 (0.21)
Average number of candidates contesting elections	9.18	8.86	0.32 (0.50)
Proportion of winners who were incumbents	0.11	0.13	−0.019 (0.049)
Average votes received by winners	48772.0	48693.7	78.3 (1621.5)
Average total votes in the constituency	128010.7	126951.2	1059.4 (3639.9)
Number of non-college educated winners in close elections	314		
Number of college educated winner in close elections	279		
Total close elections	593		

*** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Figure 5.3: Continuity of Vote Margin between Graduate and Non-Graduate (running variable)
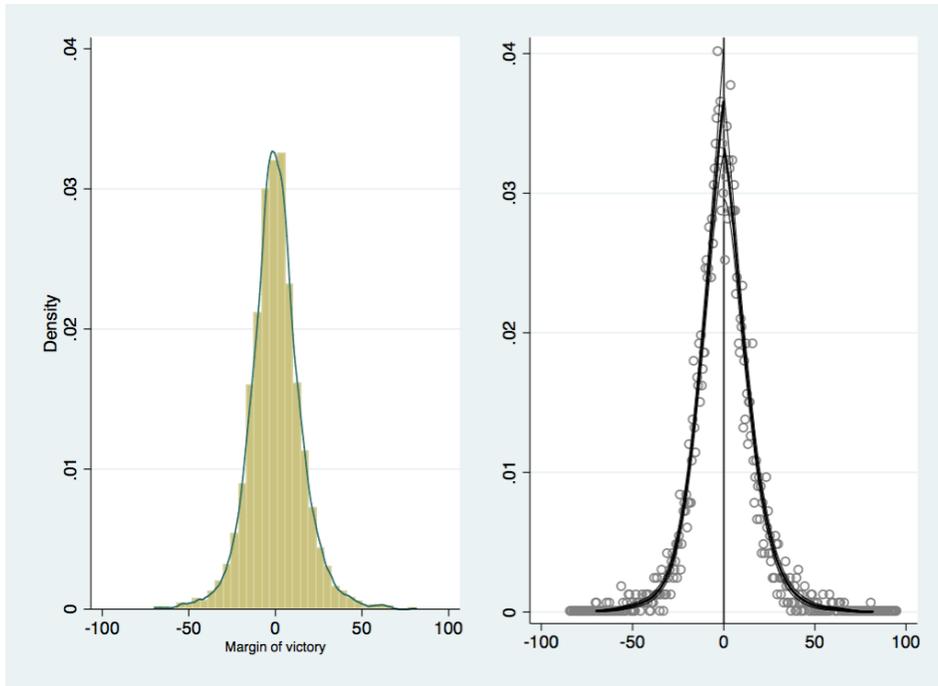
incumbency status and previous vote share are strongly correlated with winning close elections in the United States House elections, pointing to the possibility of manipulation (Caughey and Sekhon, 2011; Grimmer, Hersh, Feinstein, and Carpenter, 2011). However Eggers et al. (2015) find that such sorting in close elections is unique to U.S. House elections in the post-war period and does not hold forth in other countries including India. Along these lines, we test for differences in various constituency and candidate characteristics in close elections. Table 5.6 shows that these are not significantly different across close elections where graduate or non-graduate candidates win.

Moreover, we test for the possibility that the outcome of a close election is biased in favour of the incumbent or a party that has strategic influence over the election process. In particular, we consider whether the college graduate candidate wins in a close election between college graduate and non-college graduate candidates, and regress this outcome on the party affiliation dummies of the candidates who fought in those elections, along with other constituency and district level characteristics (column 1 of Table 5.7). In another specification we also include the history of close

elections and the proportion of college graduate winners in the district as additional explanatory variables in this regression (column 2 of Table 5.7). We find that the outcome of a close election is not significantly predicted by any of the variables related to political parties or past election outcomes.

The RD design assumptions might be violated if there is vote manipulation leading to bias in the outcome of close election. To verify that there is no maniuplation involved we check if the distribution of the vote margin is continuous around the neighborhood of zero. We plot the density of vote margin (Figure 5.3) and test if difference in the densities on either side of the zero point is significant (McCrary, 2008). The estimated difference is 0.099 and is statistically insignificant. The election commission of India which conducts elections in India is known to be independent and politically neutral and elections are considered free and fair.

In addition, to satisfy the RD assumptions one must show that districts where more graduates win in close election are not systematically different than where more non-graudates win. In Table 5.8 we compare various characteristics in the two types of districts. The districts with more graduate winners do not differ significantly in any of the characteristics from districts with more non-graduate winners in close elections.

We have shown that several of RD assumptions are valid for our data, but our RD estimates might still not have external validity. Though we cannot prove this conclusively, we do provide some indications that our results have considerable external validity.

First, we show that the identity of constituencies and districts that have close elections between graduate and non-graduates changes substantially over the elections. The average change in the proportion of a district's constituencies that experience close elections over multiple elections is 71 percent. This indicates that we are not picking up some feature of districts that have close elections and points to external validity. Also the percentange of graduate winners in close elections within a district changes by 61 percent over election cycles, indicating little correlation between graduates winning close elections and district characteristics.

Table 5.7: Probablity that College-Educated Candidate Wins in Close Elections (Dependent Variable: Dummy indicating whether winner in close election is college-educated)

	(1)	(2)
Congress parties contesting election	0.136	−0.491
	(0.207)	(0.458)
Hindu parties contesting election	0.0410	−0.0312
	(0.150)	(0.381)
Regional parties contesting election	0.185	0.136
	(0.341)	(0.404)
Left parties contesting election	−0.0726	−0.219
	(0.153)	(0.321)
Independent or other parties contesting election	−0.115	−0.256
	(0.191)	(0.472)
Reserved constituency	−0.174	−0.0310
	(0.125)	(0.322)
Proportion of urban population in district in 2004-05	−0.223	−0.0436
	(0.347)	(1.001)
Proportion of adult men who have completed primary education in 2004-05	0	0
	(.)	(.)
Proportion of adult women who have completed primary education in 2004-05	0.00308	−1.432*
	(0.490)	(0.736)
Proportion of ST population in 2004-05	−0.0426	−0.625
	(0.187)	(0.581)
Proportion of SC population in 2004-05	−0.813	0
	(0.611)	(.)
Proportion of OBC population in 2004-05	−0.0481	−1.187**
	(0.434)	(0.585)
Proportion of female population in 2004-05	0	0
	(.)	(.)
Dummy if district had close elections in past		0.373
		(0.378)
Proportion of college educated winners in past in district		0.00662
		(0.766)
Constant	0.167	1.935**
	(0.392)	(0.891)
Observations	503	236

*** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table 5.8: Comparing District Specific Characteristics across Districts-Election Years with Different Number of Educated Politician Winners in Close Elections.
Sample: District election years with close elections

	More non-college educated winners in close elections	More college educated winners in close elections	Difference
Proportion of urban population	0.22	0.20	0.02
			(0.02)
Rural adult men primary completion rate	0.66	0.66	−0.00
			(0.02)
Rural adult women primary completion rate	0.40	0.39	0.01
			(0.02)
Proportion of rural ST population	0.13	0.14	−0.01
			(0.03)
Proportion of rural SC population	0.21	0.20	0.02
			(0.01)
Proportion of rural OBC population	0.41	0.42	−0.02
			(0.03)
Proportion of rural female	0.49	0.49	−0.00
			(0.00)
Proportion of SC/ST seats	0.29	0.31	−0.02
			(0.03)
Total seats	8.00	7.71	0.30
			(0.48)
Proportion of college educated leaders win in non-close elections	0.62	0.62	−0.00
			(0.03)
Proportion non-college educated leaders win in non-close elections	0.34	0.35	−0.01
			(0.02)
Number of district-election year with more non-college educated winners in close elections	220		
Number of district-election year with more college educated winners in close elections	185		
Number of district-election year with same number of college and non-college educated winners in close elections	40		
Total number of district-election year with close elections	445		

*** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

5.5.1 Main Results: Impact of educated leaders

Learning outcomes

The second stage estimates for the effect on reading score and maths score are shown in Tables 5.9 and 5.10 respectively. In addition to the variables considered in the OLS specification, the 2SLS model includes fraction of seats that had close election between a graduate and a non-graduate in the district, third order polynomials in the vote margins in every graduate and a non-graduate election as additional covariates. We find that there is no significant effect of having a higher proportion of graduate leaders in the districts on the learning score of children in either 6–10 or 11–16 years age-group. Even the point estimates suggest that the magnitude of effect is quite low. For instance, the estimate from the full model for 6–10 year old children's reading score implies that a 10 percentage point increase in the fraction of college-educated politicians in the district would lead to a 0.0037 standard deviation increase in the reading score. The corresponding 95 percent confidence interval would be between -0.017 and 0.024. Thus, even if we take the upper bound of the confidence interval, the effect size remains negligible. Similar conclusions can be drawn from the findings on other learning outcomes including English scores presented in Table A11.

Enrollment in schools

The results for second stage of probability of being out of school using NSS data are presented in Table 5.11. We do not find any significant relationship between fraction of seats won by graduates in the district and the overall enrollment rate in the district. We repeat this analysis only for rural and urban areas separately and find no effect (not shown here). We also repeat the analysis with ASER data for rural areas with similar results (Table A7).

This is in contrast to finding by Clots-Figueras (2012) who using a similar empirical stratergy finds that female political representation decreases the chances that individual is out of school. Clots studies impact of female leaders on education and for the period 1971 to 1999, as compared to our study which studies impact of formal education of the leader for more recent period 2004-2014.

We also consider grade progression as another outcome variable. [18] Unlike current enrollment,

---

[18] We construct grade progression as the ratio of actual grade attained to ideal grade for age; hence it is expressed as: $Grade/(Age - 6)$. The numerator measures the actual grade attained, which is the last grade if a child has dropped

Table 5.9: 2SLS Estimates of the Effect on Children's Reading Score

	Reading Score					
	6–10 age-group			11–16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won by a graduate	-0.014	0.086	0.037	-0.020	0.085	0.134
	(0.123)	(0.124)	(0.104)	(0.096)	(0.098)	(0.103)
Close election fraction	0.021	0.000	-0.025	-0.069	-0.128**	-0.150**
	(0.069)	(0.065)	(0.053)	(0.048)	(0.055)	(0.060)
Female child		0.007*	0.008*		-0.009**	-0.008*
		(0.004)	(0.004)		(0.005)	(0.005)
Child's mother went to school		0.258***	0.258***		0.222***	0.221***
		(0.006)	(0.006)		(0.006)	(0.006)
Household size		-0.037***	-0.036***		-0.030***	-0.030***
		(0.002)	(0.002)		(0.002)	(0.002)
Square of household size		0.002***	0.002***		0.001***	0.001***
		(0.000)	(0.000)		(0.000)	(0.000)
House semi-pucca		0.093***	0.090***		0.093***	0.091***
		(0.005)	(0.005)		(0.006)	(0.006)
House pucca		0.206***	0.199***		0.161***	0.158***
		(0.007)	(0.006)		(0.007)	(0.007)
Household owns television		0.141***	0.138***		0.103***	0.102***
		(0.005)	(0.005)		(0.005)	(0.004)
Household owns mobile phone		0.120***	0.129***		0.144***	0.150***
		(0.005)	(0.004)		(0.006)	(0.005)
Household's electricity use observed		0.030***	0.033***		0.034***	0.036***
		(0.005)	(0.005)		(0.005)	(0.004)
Village has electricity		0.030**	0.046***		0.025**	0.034***
		(0.012)	(0.012)		(0.013)	(0.013)
Village has pucca road		0.027***	0.028***		0.033***	0.033***
		(0.006)	(0.006)		(0.005)	(0.005)
Village has ration shop		0.025***	0.025***		0.021***	0.022***
		(0.006)	(0.006)		(0.005)	(0.005)
Village has bank		0.044***	0.042***		0.013***	0.013***
		(0.005)	(0.005)		(0.004)	(0.004)
Vote margins: third order polynomial	Yes	Yes	Yes	Yes	Yes	Yes
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Age cohort fixed effects		Yes	Yes		Yes	Yes
Child, household & village controls		Yes	Yes		Yes	Yes
District controls			Yes			Yes
State by year fixed effects			Yes			Yes
Observations	1,262,927	896,820	869,544	1,272,189	928,499	904,149
Number of districts 563	563	545	563	563	545	
First Stage F-stat	53.68	60.36	55.37	60.98	66.03	60.88

Notes: Robust standard errors clustered at the district level are in parentheses. Close elections are defined as between a graduate and a non-graduate in which the difference in vote share between the winner and the runner up is less than 3 percent. District controls are baseline characteristics interacted with year dummies; they include urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas estimated from 2004-05 National Sample Survey data. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table 5.10: 2SLS Estimates of the Effect on Children's Mathematics Score

	Mathematics Score					
	6–10 age-group			11–16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won by a graduate	-0.115	0.037	0.014	-0.164	0.007	0.108
	(0.150)	(0.155)	(0.112)	(0.168)	(0.179)	(0.131)
Close election fraction	0.009	-0.057	-0.082	-0.066	-0.141	-0.163**
	(0.077)	(0.077)	(0.055)	(0.078)	(0.091)	(0.063)
Female child		-0.043***	-0.043***		-0.084***	-0.082***
		(0.004)	(0.004)		(0.006)	(0.006)
Child's mother went to school		0.251***	0.253***		0.258***	0.259***
		(0.006)	(0.006)		(0.005)	(0.005)
Household size		-0.036***	-0.035***		-0.036***	-0.035***
		(0.002)	(0.002)		(0.002)	(0.002)
Square of household size		0.002***	0.001***		0.002***	0.002***
		(0.000)	(0.000)		(0.000)	(0.000)
House semi-pucca		0.084***	0.084***		0.095***	0.098***
		(0.005)	(0.005)		(0.006)	(0.006)
House pucca		0.209***	0.202***		0.196***	0.190***
		(0.006)	(0.006)		(0.006)	(0.006)
Household owns television		0.139***	0.138***		0.126***	0.128***
		(0.005)	(0.005)		(0.004)	(0.004)
Household owns mobile phone		0.116***	0.124***		0.127***	0.137***
		(0.004)	(0.004)		(0.005)	(0.005)
Household's electricity use observed		0.036***	0.035***		0.040***	0.041***
		(0.005)	(0.005)		(0.005)	(0.004)
Village has electricity		0.033***	0.042***		0.026*	0.019
		(0.012)	(0.013)		(0.014)	(0.013)
Village has pucca road		0.026***	0.027***		0.027***	0.029***
		(0.005)	(0.005)		(0.005)	(0.005)
Village has ration shop		0.027***	0.029***		0.020***	0.025***
		(0.006)	(0.006)		(0.005)	(0.005)
Village has bank		0.045***	0.043***		0.018***	0.016***
		(0.005)	(0.005)		(0.005)	(0.004)
Vote margins: third order polynomial	Yes	Yes	Yes	Yes	Yes	Yes
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Age cohort fixed effects		Yes	Yes		Yes	Yes
Child, household & village controls		Yes	Yes		Yes	Yes
District controls			Yes			Yes
State by year fixed effects			Yes			Yes
Observations	1,251,958	892,605	865,509	1,267,391	926,214	901,972
Number of districts	563	563	545	563	563	545
First Stage F-stat	53.68	60.30	55.25	60.94	65.93	60.77

Notes: Robust standard errors clustered at the district level are in parentheses. Close elections are defined as between a graduate and a non-graduate in which the difference in vote share between the winner and the runner up is less than 3 percent. District controls are baseline characteristics interacted with year dummies; they include urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas estimated from 2004-05 National Sample Survey data. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table 5.11: 2SLS Estimates of the Impact on Proportion Out of School (NSS Data)

| | Proportion out of school | | | | | |
| | 6–10 age-group | | | 11–16 age-group | | |
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won by a graduate	0.032	0.028	0.021	0.055	0.056	0.045
	(0.043)	(0.043)	(0.038)	(0.038)	(0.038)	(0.036)
Vote margins: third order polynomial	Yes	Yes	Yes	Yes	Yes	Yes
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Age cohort fixed effects		Yes	Yes		Yes	Yes
Child, household controls		Yes	Yes		Yes	Yes
District controls			Yes			Yes
State by year fixed effects			Yes			Yes
Observations	108,521	108,493	108,493	131,693	131,651	131,651
Number of districts	574	574	574	574	574	574
First Stage F-stat	75.65	75.75	75.52	80.17	80.27	82.38

Notes: Robust standard errors clustered at the district level are in parentheses. Close elections are defined as between a graduate and a non-graduate in which the difference in vote share between the winner and the runner up is less than 3 percent. District controls are baseline characteristics interacted with year dummies; they include urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas estimated from 2004-05 National Sample Survey data. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

attainment of grade commensurate with one's age, or grade progression, is sensitive to the age when a child dropped out. School dropout is still a matter of concern especially at the secondary level of education when opportunity cost of education increases for boys who start entering the labour market, and girls are often married off or engaged in household work. For children who are currently enrolled, this variable would capture various other aspects such as whether the child had late school-starting age or went through grade rentention – which are indicated as important determinants of human capital in the literature. Even when we consider grade progression as an alternative dependent variable, we do not find any significant effect of graduate versus non-graduate leaders on this outcome (Table A9).

out, and zero if a child has never been enrolled in school. The denominator captures the ideal grade for age under the assumption that a child is supposed to enter the first grade in school at the age of six years. Thereafter, the ideal grade attainment increases by one for every incremental year.

School Infrastructure and Funding

We analyze the impact of having educated representatives on a range of school indicators including school physical infrastructure, financial grants and other outcomes using smaller sample of data on schools collected by ASER (children's and teachers' attendance and midday meal program in schools). These outcomes could be more directly impacted by leaders as they have a more direct role to play in deciding and channeling funding to schools. DISE collects data from all recognized schools in India (> 1 million each year) and we use district averages for outcomes in the regression analysis. ASER collects school data from one government school in each rural village[19] it surveys (56-76K across the years). We do separate analysis using both the datasets.

We find no significant impact of having higher proportion of graduate leaders on school infrastructure using both DISE (Table 5.12) and ASER data[20] (Table A14). Similarly, we find no evidence that graduate politicians result in any change in the extent of grants received by schools using either DISE (Table 5.12) or ASER[21] data (Table A14). We also do not find any impact on prevalence of midday meals in schools using data from ASER.

Finally, we analyze if graduate politicians improve teachers' and children's attendance in schools[22] using data from ASER. For children's attendance we find a small negative but insignificant impact of having an educated politician, but impact on teacher's attendance is positive and insignificant (Table A13).

Overall, educated political representatives do not make significant difference across a range of schooling outcomes. This finding is in line with our other results that educated leaders are no better at improving learning outcomes and out-of-school status than non-graduate leaders. Even if one considers learning outcomes and enrollment as a difficult outcome to be impacted by political leaders, through funding and lobbying politicians can impact schooling facilities directly. But we do not find consistent impact of educated politicians on schooling services for their constituents.

---

[19]If the village has a government school for class 1 to 7/8 that is preferred over other schools to be visited. If village does not have school for classes 1 to 7/8 then the primary school (teaching classes 1 to 4/5) with highest enrollment is visited.

[20]PCA for school infrastructure for ASER data is created using dummies for presence of blackboard and learning material in class 2 and class 4, drinking water, toilet and library books in school.

[21]For ASER we combine three grants (school development, school maintenance and teacher learning material) to generate a PCA which is used as the dependent variable.

[22]Children and teachers attendance on the day of the survey is observed by the interviewer and recorded. The attendance ratio is defined as the ratio of total children attending the school on the day of the survey (teachers present) to the total number of enrolled students based on the register (appointed teachers).

Table 5.12: 2SLS Estimates of the Impact on Schools (DISE data)

	PCA for school facilites		Average school grants	
	(1)	(2)	(3)	(4)
Fraction of	0.179	-0.134	1337.4	4412.8
seats won by a graduate	(0.240)	(0.181)	(7041.3)	(6343.2)
Vote margins: third order polynomial	Yes	Yes	Yes	Yes
District fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
District controls		Yes		Yes
State by year fixed effects		Yes		Yes
Number of districts	2607	2167	2692	2240
First stage F statistics	60	57	60	55

Robust standard errors clustered at the district level are in parentheses. Close elections are defined as between a graduate and a non-graduate in which the winner beat the runner up by less than 3 percent of votes. District controls are interacted with year dummies and include data from 2004-05 National Sample Survey on urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

5.5.2 Robustness Analysis

We put our results to robustness checks in this section to investigate if varying the specification or range of political leaders who are supposed to have an impact changes our results.

Different Education Cuttoff

In our baseline analysis we compared college graduate leader's achievements in education sector with non-graduate leaders. It might be the case that the our results depend on the education cuttoffs we have chosen and that other comparisons would yield different results. We test robustness of our results to this by altering the leader's education cuttoffs to higher secondary (12th grade) and secondary education (10th grade). Leaders who have completed higher secondary do not perform any better at providing education outcomes for their constituents than leaders who have not completed higher secondary. And similarly leaders who have completed secondary education do not perform any better than leaders who have not completed secondary education in any of the outcomes we investigate (Table 5.13).

Table 5.13: Robustness: Alternative Cutoff of Leader's Education

Panel B: Impact on reading scores

	6-10 age-group		11-16 age-group	
	(1)	(2)	(3)	(4)
Fraction of seats won by 10th pass	0.299 (0.228)		0.204 (0.189)	
Fraction of seats won by 12th pass		-0.002 (0.104)		0.062 (0.093)
Observations	869,544	869,544	904,149	904,149
Number of districts	545	545	545	545
First stage F statistics	56.58 10th	46.14 12th	50.69 10th	50.21 12th
Leaders education	Complete	Complete	Complete	Complete

Panel B: Impact on mathematics scores

	6-10 age-group		11-16 age-group	
	(1)	(2)	(3)	(4)
Fraction of seats won by 10th pass	0.199 (0.176)		0.257 (0.244)	
Fraction of seats won by 12th pass		0.031 (0.103)		0.064 (0.113)
Observations	865,509	865,509	901,972	901,972
Number of districts	545	545	545	545
First stage F statistics	56.51 10th	46.24 12th	50.66 10th	50.23 12th
Leaders education	Complete	Complete	Complete	Complete

Panel C: Impact on being out of school (NSS data)

	6-10 age-group		11-16 age-group	
	(1)	(2)	(3)	(4)
Proportion of seats won by 10th pass	–0.0559 (0.0796)		–0.101 (0.0985)	
Proportion of seats won by 12th pass		–0.0204 (0.0441)		0.0189 (0.0488)
Observations	107,391	107,391	130,354	130,354
Number of districts	567	567	567	567
First stage F statistics	25 10th	54 12th	23 10th	52 12th
Leaders education	Complete	Complete	Complete	Complete

Panel D: Impact on schools (DISE data)

	PCA for school facilites		Average school grants	
	(1)	(2)	(3)	(4)
Fraction of seats won by 10th pass	0.557 (0.465)		6266.0 (14984.8)	
Fraction of seats won by 12th pass		0.155 (0.237)		1605.8 (5225.5)
Observations	2,378	2,378	2,453	2,453
Number of districts	518	518	519	519
First stage F statistics	30 10th	44 12th	30 10th	45 12th
Leaders education	Complete	Complete	Complete	Complete

168

Note: The regressions include the full set of controls as reported in Table 5.9. Robust standard errors clustered at the district level are in parentheses. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Different Lag Periods

It might be the case that the policies an educated leader implements may take time to show an effect on the outcomes. A priori it is not clear how long it takes for a policy to show effect after its implementation, therefore this remains an empirical issue. The main results discussed above consider the average education level (whether college graduate or not) of leaders in the district over the last three years (current year and the past two years). This measure is similar to the one used by Clots-Figueras (2012) who finds a significant positive effect of having female versus male politicians on education of citizens. As a robustness exercise, we use different lag periods, in particular the average over the last four years (Lag 3) and two years (Lag 1), to measure the explanatory variable. Irrespective of the lag period used, the effect of college graduate leaders, in comparison with non-college graduate leaders, remains statistically insignificant for all outcomes (Table 5.14).

Alternative Definition of Close Election

We define close elections as those where the margin of victory is less than 3 percent. We test whether the results are robust to alternative cut-off points of 1, 2, and 4 percent levels, and find that the results remain qualitatively unchanged (Table 5.15).

Varying Degrees of Polynomials in Vote Margins

Given the fuzzy regression discontinuity design of our empirical model, it is important to control for polynomials in the vote margins so that any effect of the vote margin itself is controlled in a flexible way. This ensures that the instrument exploits variation only in the close neighborhood around the discontinuity to justify a quasi-random assignment of treatment. While the main regressions include third order polynomials, our results (not shown here) are robust to using only first and second order polynomials as well.

Different Measures of Learning Outcomes

We have so far used the learning outcomes as continuous variable, but learning outcome is an ordinal variable. Estimating the regressions using ordered probit or ordered logit would be more

Table 5.14: Robustness: Different Lag Periods

Panel A: Impact on reading scores

	6–10 age-group		11–16 age-group	
	Lag 1	Lag 3	Lag1	Lag 3
Fraction of seats won by a graduate	-0.003	0.002	0.103	0.117
	(0.089)	(0.148)	(0.085)	(0.127)
Observations	998,524	712,455	1,031,315	746,926
Number of districts	549	543	549	543
First Stage F-stat	55.48	50.67	62.35	53.93

Panel B: Impact on mathematics scores

	6–10 age-group		11–16 age-group	
	Lag 1	Lag 3	Lag1	Lag 3
Fraction of seats won by a graduate	-0.026	-0.018	0.108	0.046
	(0.095)	(0.154)	(0.114)	(0.161)
Observations	993,964	709,434	1,028,937	745,292
Number of districts	549	543	549	543
First Stage F-stat	55.50	50.41	62.25	53.80

Panel C: Impact on being out of school (NSS Data)

	6–10 age-group		11–16 age-group	
	Lag 1	Lag 3	Lag1	Lag 3
Fraction of seats won by a graduate	-0.009	0.054	0.019	0.043
	(0.029)	(0.046)	(0.029)	(0.042)
Observations	125,459	82,413	150,035	100,711
Number of districts	578	573	578	573
First Stage F-stat	68.72	56.16	76.76	60.98

Panel D: Impact on schools (DISE data)

	PCA for school facilites		Average school grants	
	Lag 1	Lag 3	Lag1	Lag 3
Fraction of seats won by a graduate	-0.130	-0.0915	5913.3	4875.2
	(0.171)	(0.213)	(5967.6)	(6980.6)
Observations	2483	2181	2581	2248
Number of districts	520	518	521	519
First stage F statistics	49	49	49	46

Note: The regressions include the full set of controls as reported in Table 5.9. Robust standard errors clustered at the district level are in parentheses. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level. 'Lag 1' considers the average education of leaders in the district over the past 1 year and the current year (i.e., average over last 2 years). 'Lag 3' considers the average education of leaders over the past 3 years and the current year (i.e. average over last 4 years).

Table 5.15: Robustness: Alternative Definitions of Close Election Margin

Panel A: Impact on reading scores

	6–10 age-group			11–16 age-group		
	Close election margin			Close election margin		
	1 %	2 %	4 %	1 %	2 %	4 %
Fraction of seats won by a graduate	-0.041	0.069	0.085	0.154	0.110	0.138
	(0.182)	(0.106)	(0.091)	(0.129)	(0.094)	(0.085)
Observations	869,544	869,544	869,544	904,149	904,149	904,149
Number of districts	545	545	545	545	545	545
First Stage F-stat	43.02	83.66	75.16	43.43	88.94	79.56

Panel B: Impact on mathematics scores

	6–10 age-group			11–16 age-group		
	Close election margin			Close election margin		
	1 %	2 %	4 %	1 %	2 %	4 %
Fraction of seats won by a graduate	-0.022	0.101	0.043	0.126	0.166	0.097
	(0.181)	(0.109)	(0.095)	(0.191)	(0.124)	(0.107)
Observations	865,509	865,509	865,509	901,972	901,972	901,972
Number of districts	545	545	545	545	545	545
First Stage F-stat	42.82	83.25	74.96	43.28	88.68	79.44

Panel C: Impact on being out of school (NSS data)

	6–10 age-group			11–16 age-group		
	Close election margin			Close election margin		
	1 %	2 %	4 %	1 %	2 %	4 %
Fraction of seats won by a graduate	0.074	0.049	0.037	0.007	0.068*	0.024
	(0.056)	(0.039)	(0.034)	(0.047)	(0.035)	(0.033)
Observations	108,493	108,493	108,493	131,651	131,651	131,651
Number of districts	574	574	574	574	574	574
First Stage F-stat	51.21	81.26	94.63	50.27	85.18	103.9

Panel D: Impact on schools (DISE data)

	PCA for school facilites			Average school grants		
	Close election margin			Close election margin		
	1 %	2 %	4 %	1 %	2 %	4 %
Fraction of seats won by a graduate	0.385	0.0697	-0.0108	2271.7	2195.2	-807.2
	(0.297)	(0.196)	(0.170)	(7882.2)	(6926.9)	(6032.8)
Observations	2378	2378	2378	2453	2453	2453
Number of districts	518	518	518	519	519	519
First stage F statistics	48	74	69	46	72	69

Note: The regressions include the full set of controls as reported in Table 5.9. Robust standard errors clustered at the district level are in parentheses. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

appropriate but they limit the estimation models we can use. To test the robustness of our results to ordinal nature of the dependent variable we use binary variables indicating the level of learning as alternative dependent variable. For reading, these variables reflect whether a child can read letters, words, short paragraph and short stories. For mathematics, each of them measure one of the categories from single-digit number recognition, double-digit number recognition, two-digit subtraction with carry over, and three digit by one digit division. Linear probability models are estimated using 2SLS method following the same specification as the main regressions presented above. Table 5.16 shows that there is no differential impact of having a college graduate leader over a non-college graduate leader on any of these outcomes.

5.5.3 Heterogeneity Analysis

We explore if the effect of educated politician is heterogeneous with respect to various district, household and child specific characteristics. Although we do not find any significant effect in the overall sample, it is possible that college graduate leaders are instrumental in implementing policies that benefit certain sections of the society. For instance, if educated leaders are motivated towards reducing inequality, it may benefit poorer regions. Such effects may not show up in the aggregate sample. Therefore, dividing the sample into subsamples, the 2SLS model with full set of control variables is estimated for each subsample.

Table 5.16: Robustness: Binary Indicators for Different Levels of Learning as Outcome Variable

Panel A: Impact on various levels of reading skill

	6–10 age-group				11–16 age-group			
	Letter	Word	Paragraph	Story	Letter	Word	Paragraph	Story
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fraction of seats won by a graduate	0.006	0.006	0.007	0.016	0.039	0.048	0.049	0.046
	(0.042)	(0.044)	(0.046)	(0.033)	(0.041)	(0.046)	(0.051)	(0.054)
Observations	951,476	951,476	951,476	951,476	1,016,266	1,016,266	1,016,266	1,016,266
Number of districts	545	545	545	545	545	545	545	545
First Stage F-stat	52.77	52.77	52.77	52.77	59.41	59.41	59.41	59.41

Panel B: Impact on various levels of math skill

	6–10 age-group				11–16 age-group			
	1 digit	2 digit	Subtraction	Division	1 digit	2 digit	Subtraction	Division
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fraction of seats won by a graduate	0.025	-0.024	0.008	-0.013	0.035	0.038	0.036	0.023
	(0.042)	(0.051)	(0.047)	(0.025)	(0.042)	(0.047)	(0.062)	(0.069)
Observations	951,476	951,476	951,476	951,476	1,016,266	1,016,266	1,016,266	1,016,266
Number of districts	545	545	545	545	545	545	545	545
First Stage F-stat	52.77	52.77	52.77	52.77	59.41	59.41	59.41	59.41

Note: The regressions include the full set of controls as reported in Table 5.9. Robust standard errors clustered at the district level are in parentheses. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level. The dependent variables are binary indicators, hence the regressions are linear probability models. For reading skill, the different levels are reading letters, words, a short paragraph (a class 1 level text), and a short story (a class 2 level text). For math skill, the different levels are single-digit number recognition, double-digit number recognition, two-digit subtraction with carry over, and three digit by one digit division (corresponding to what students are expected to know in grade 3 or 4).

173

Heterogeneity in poverty, urbanization and education levels in districts

The first analysis in this section tests if the effect varies with the baseline poverty rates in the districts. Based on estimates from the NSS 2004-05 data, we categorize the districts into high, medium and low poverty rate groups. High poverty districts are ones where more than 67 percent of rural population has consumption expenditure below the 40th percentile of the nation-wide distribution of consumption expenditure. Low poverty districts have less than 31 percent of the rural population having consumption expenditure below this level. Medium poverty districts are the ones that have 31-67 percent of rural population in this category. The results are presented in Table 5.17. Having college graduate leaders does not make any significant difference in math score irrespective of the district poverty rates. For reading score, we find significant positive effect at 10 percent levels for the 11–16 year old children residing in the high poverty districts. Probability of being out of school also declines significantly for low poverty districts, but school infrastructure is worse in low poverty districts if they have more graduate leaders. For all the other cases the effect is insignificant, we do not find any compelling evidence that college graduate leaders may be effective in either rich or poor districts.

Next we test if the impact of educated leaders on various outcomes differs by baseline level of urbanization and initial level of education in the districts. Districts with urbanization rate of 15 percent or less are categorized as low urbanization districts, while others are categorized as high urbanization. We also group districts below and above median level of primary education completion rates of adults in 2004-05. From Table 5.18, we find that there except for a few cases for low education level districts, there is no significant effect of educated leaders across rural/urban or different urbanization or education levels of districts. There is negative significant effect of graduate leaders in districts with baseline low education levels on learning scores (both reading and mathematics) for 6–10 age group and enrollment in 11–16 age group.

Heterogeneity in household economic status

Heterogeneity based on households economic status are explored next. From ASER data we use the type of material used for construction of house based as proxy for economic status of the household. We classify households into Kutcha House (houses made out of mud, thatch or low quality materials), Semi Pucca house (houses made of mix of low quality and more permanent material) and Pucca house (houses made of high quality material). We present the results for these

Table 5.17: Heterogeneity in Impact - Poverty Level in Districts

Panel A: Impact on reading scores

	6-10 age-group			11-16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won	0.089	0.074	0.206	0.406*	0.080	0.074
by graduates	(0.247)	(0.143)	(0.168)	(0.244)	(0.138)	(0.179)
Observations	224,090	430,746	214,708	237,137	440,698	226,314
Number of districts	143	267	135	143	267	135
First stage F statistics	14	33	73	15	36	72
Districts in sample	High Poverty	Medium Poverty	Low Poverty	High Poverty	Medium Poverty	Low Poverty

Panel B: Impact on mathematics scores

	6-10 age-group			11-16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won	–0.135	0.118	0.202	–0.044	0.190	–0.084
by graduates	(0.246)	(0.145)	(0.219)	(0.284)	(0.151)	(0.265)
Observations	223,166	428,671	213,672	236,598	439,645	225,729
Number of districts	143	267	135	143	267	135
First stage F statistics	14	33	73	15	36	72
Districts in sample	High Poverty	Medium Poverty	Low Poverty	High Poverty	Medium Poverty	Low Poverty

Panel C: Impact on proportion out of school (NSS data)

	6-10 age-group			11-16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won	0.050	0.001	–0.147***	0.025	–0.012	–0.022
by graduates	(0.064)	(0.057)	(0.038)	(0.112)	(0.060)	(0.062)
Observations	22,714	51,515	23,070	28,094	60,858	29,128
Number of districts	131	252	127	131	252	127
First stage F statistics	16	30	46	17	36	51
Districts in sample	High Poverty	Medium Poverty	Low Poverty	High Poverty	Medium Poverty	Low Poverty

Panel D: Impact on schools (DISE data)

	PCA for school facilites			Average school grants		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won	–0.0676	0.254	–0.857*	283.6	7960.7	2492.5
by graduates	(0.641)	(0.278)	(0.369)	(13109.4)	(6024.5)	(9855.8)
Observations	607	1,163	608	624	1,193	636
Number of districts	132	253	133	132	253	134
First stage F statistics	6	33	39	6	31	38
Districts in sample	High Poverty	Medium Poverty	Low Poverty	High Poverty	Medium Poverty	Low Poverty

Note: See notes to Table 5.9. All regressions include the full set of control variables as reported in Table 5.9. High poverty districts are ones where more than 67% of rural population is in the bottom 40th national consumption percentile. Rich poverty districts have less than 31% of the rural population in the bottom 40th percentile by consumption. Medium poverty districts are ones with between 31-67% of rural population in the bottom 40th consumption percentile.

175

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the district level are in parentheses.

Table 5.18: Heterogenity in Impact - Urbanization and Level of Education in District

Panel A: Impact on reading scores

	6-10 age-group				11-16 age-group			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fraction of seats won	0.004	0.100	–0.244**	0.279	0.066	0.249	–0.020	0.278
by graduates	(0.099)	(0.181)	(0.118)	(0.210)	(0.116)	(0.155)	(0.100)	(0.193)
Observations	421,583	447,961	432,794	436,750	416,691	487,458	429,495	474,654
Number of districts	258	287	271	274	258	287	271	274
First stage F statistics	44	34	29	31	50	31	32	28
Districts in sample	Low urb.	High urb.	Low ed. level	High ed. level	Low urb.	High urb.	Low ed. level	High ed. level

Panel B: Impact on mathematics scores

	6-10 age-group				11-16 age-group			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fraction of seats won	0.059	–0.049	–0.246*	0.334	0.103	0.147	–0.062	0.344
by graduates	(0.108)	(0.207)	(0.128)	(0.219)	(0.150)	(0.223)	(0.139)	(0.241)
Observations	419,532	445,977	430,978	434,531	415,631	486,341	428,486	473,486
Number of districts	258	287	271	274	258	287	271	274
First stage F statistics	44	34	29	31	50	31	32	28
Districts in sample	Low urb.	High urb.	Low ed. level	High ed. level	Low urb.	High urb.	Low ed. level	High ed. level

Panel C: Impact on proportion out of school (NSS data)

	6-10 age-group				11-16 age-group			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fraction of seats won	0.034	0.024	0.015	–0.069	0.072	–0.007	0.119*	–0.024
by graduates	(0.060)	(0.035)	(0.047)	(0.068)	(0.067)	(0.057)	(0.067)	(0.060)
Observations	43,333	53,966	46,645	50,654	49,282	68,798	53,888	64,192
Number of districts	239	271	254	256	239	271	254	256
First stage F statistics	43	38	36	37	47	38	41	34
Districts in sample	Rural	Urban	Low ed. level	High ed. level	Rural	Urban	Low ed. level	High ed. level

Panel D: Impact on schools (DISE data)

	PCA for school facilites				Average school grants			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fraction of seats won	–0.0887	–0.365	–0.243	–0.0197	5100	–6018.5	–540.9	14588.8
by graduates	(0.178)	(0.398)	(0.220)	(0.328)	(6870.9)	(3935.4)	(5118.0)	(15428.6)
Observations	2,378	2,335	1,168	1,210	2,453	2,419	1,183	1,270
Number of districts	518	510	257	261	519	514	257	262
First stage F statistics	54	57	28	31	51	54	26	30
Districts in sample	Rural	Urban	Low ed. level	High ed. level	Rural	Urban	Low ed. level	High ed. level

Note: See notes to Table 5.9. All regressions include the full set of control variables as reported in Table 5.9. Districts with urbanization rate of 15 percent or less are categorized as low urbanization districts, while others are categorized as high urbanization. Districts where 51 or lower percentage of adults have completed primary education are low education level districts and other districts are classified as high education level districts. We use NSS data for proportion of children out of school in both rural and urban areas and ASER for learning levels in rural areas.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the district level are in parentheses.

Table 5.19: Heterogenity in Impact - Type of Residence

Panel A: Impact on reading scores

	6-10 age-group			11-16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won	–0.034	–0.027	0.068	0.223	0.155	–0.026
by graduates	(0.148)	(0.134)	(0.096)	(0.173)	(0.142)	(0.072)
Observations	280,040	268,224	321,280	270,405	277,659	356,085
Number of districts	545	545	545	545	545	545
First stage F statistics	36	52	60	39	58	63
Households in sample	Kutcha House	Semi Pucca House	Pucca House	Kutcha House	Semi Pucca House	Pucca House

Panel B: Impact on mathematics scores

	6-10 age-group			11-16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won	–0.026	0.026	–0.004	0.201	0.195	–0.056
by graduates	(0.164)	(0.142)	(0.110)	(0.206)	(0.177)	(0.105)
Observations	278,575	266,964	319,970	269,624	276,995	355,353
Number of districts	545	545	545	545	545	545
First stage F statistics	36	52	60	39	58	63
Households in sample	Kutcha House	Semi Pucca House	Pucca House	Kutcha House	Semi Pucca House	Pucca House

Panel C: Impact on proportion out of school (NSS data)

	6-10 age-group			11-16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won	0.004	–0.061*	–0.082*	0.011	0.062	0.132***
by graduates	(0.064)	(0.035)	(0.043)	(0.080)	(0.057)	(0.049)
Observations	40,411	38,115	18,766	41,251	47,878	28,946
Number of districts	508	510	506	508	510	507
First stage F statistics	56	67	60	64	75	70
Households in sample	Bottom 30 Perc	Middle 40 Perc	Top 30 Perc	Bottom 30 Perc	Middle 40 Perc	Top 30 Perc

Note: See notes to Table 5.9. All regressions include the full set of control variables as reported in Table 5.9. A pucca house is one whose roof and walls are made of permanent material (brick, cement etc). A kutcha house is one whose roof and flooring are made of temporary material (mud, grass etc). A semi-pucca house is one that has fixed walls but roof is made of temporary material.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the district level are in parentheses.

Table 5.20: Heterogenity in Impact - Gender and Mothers Education

Panel A: Impact on reading scores

	6-10 age-group				11-16 age-group			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fraction of seats won by graduates	–0.008 (0.104)	0.083 (0.114)	–0.099 (0.132)	0.088 (0.108)	0.156 (0.101)	0.116 (0.116)	0.084 (0.140)	0.117 (0.095)
Observations	462,741	406,803	363,983	505,561	471,751	432,398	416,941	487,208
Number of districts	545	545	545	545	545	545	545	545
First stage F statistics	55	56	41	66	61	62	49	66
Sample	Boys	Girls	Mother did not go to school	Mother went to school	Boys	Girls	Mother did not go to school	Mother went to school

Panel B: Impact on mathematics scores

	6-10 age-group				11-16 age-group			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fraction of seats won by graduates	0.020 (0.110)	0.001 (0.125)	–0.152 (0.144)	0.093 (0.114)	0.178 (0.121)	0.036 (0.152)	–0.009 (0.161)	0.139 (0.124)
Observations	460,608	404,901	362,266	503,243	470,663	431,309	415,837	486,135
Number of districts	545	545	545	545	545	545	545	545
First stage F statistics	55	56	41	66	61	62	49	66
Sample	Boys	Girls	Mother did not go to school	Mother went to school	Boys	Girls	Mother did not go to school	Mother went to school

Panel C: Impact on proportion out of school (NSS data)

	6-10 age-group				11-16 age-group			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fraction of seats won by graduates	–0.042 (0.039)	0.029 (0.043)	–0.029 (0.075)	–0.034 (0.029)	–0.022 (0.041)	0.139** (0.062)	0.133** (0.068)	–0.005 (0.034)
Observations	52,173	45,126	30,810	42,839	63,329	54,751	43,895	55,050
Number of districts	510	510	495	510	510	510	505	510
First stage F statistics	73	71	55	84	79	79	62	87
Sample	Boys	Girls	Mother did not go to school	Mother went to school	Boys	Girls	Mother did not go to school	Mother went to school

Note: See notes to Table 5.9. All regressions include the full set of control variables as reported in Table 5.9.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors clustered at the district level are in parentheses.

sub-samples in Table 5.19. College graduate politicians perform the same as non-college graduate leaders for all kinds of residence.

For NSS data we have information on consumption of households and we use per-capita consumption to divide households into three groups (bottom 30 percent, middle 40 percent and top 30 percent). We find that for 6-10 age group there is a positive impact significant at 10 percent level of significance for the richer households, but negative impact on enrollment for 11-16 age group for the richest households.

Heterogeneity in Individual Characteristics

Table 5.20 presents results of heterogeneity analysis based on child level characteristics: gender and whether mother went to school. If graduate politicians are more aware of gender disparities in society and want to elevate those then we can expect to see bigger impact in education outcomes on girls than boys. Mother's education level is an important indicator of the household environment and has direct bearing on child's learning outcomes. Children with mothers who have not gone to school tend to have lower schooling outcomes and any improvement in schooling environment might have a differential impact on them as compared to other children. College graduate political leaders do not perform better than other politicians across these groupings for learning outcomes. But graduate leaders led to lower enrollment rates for girls and children with mothers who have not gone to school in 11-16 age group.[23]

5.6 Conclusion

Formal education of the leader has been used as a proxy for quality of the leader without much evidence on the efficacy of educated leaders. Recent changes in laws in some states in India have used this notion to bar less educated citizens from contesting elections. In this paper, we investigate whether having more-educated state representatives results in better learning and schooling outcomes. We find that more-educated politicians are no better at improving elementary education for their constituents than less-educated politicians. Learning outcomes, enrollment, attendance, school infrastructure and funding are not significantly impacted if the elected representative is

---

[23]Since DISE data on schools does not have household or individual level characteristics, those are not included in this analysis.

more educated. In the field of elementary education leaders with more formal education are not found to be more competent or effective.

Our results are subject to several important caveats. First, the results are specific to the outcomes in education. Even though education is an important development outcome that ought to be a priority of any good leader, it might be the case the educated leaders have other priorities which they focus on. Educated leaders might get better outcomes than less educated leaders in other areas such as economic growth, health or environment which are not analyzed in this paper. Also even though state leaders have a role to play in education policy and implementation, yet it might be a hard outcome for leaders to influence. Hence we test the impact of educated leader on a range of education outcomes including school level outcomes which are arguably more direct and easier to influence for a policy-maker. Moreover, we link our results to the existing literature which has found impact of state leaders on education, although we cannot completely rule out the possibility that education is an outcome on which state leaders may not have an impact.

Second, we only analyze state leaders and the impact of education of leaders might be different at the national or the local level. National and local leaders might have different set of responsibilities and roles than state level leaders and it might be the case that education plays a significant role for them but not for state leaders.

Third, due to limitations of data we analyze leader characteristics only at the district level and not at individual constituency level for which they are elected. Leaders are elected for constituency and each district has on average nine constituencies. But the lowest unit of identification in the data on outcomes is the district, so we aggregate leader characteristics at the district level before ascertaining the impact. Our results might be sensitive to this averaging of leader characteristics at the district level.

Fourth, our identification strategy uses close elections to estimate the impact of educated leaders. Even though we control for the proportion of close elections in the district and provide some suggestive evidence of external validity of our results, leaders (educated and less-educated) might behave differently in close elections vis-a-vis non-close elections. Maybe in non-close elections educated leaders face less political competition and spend resources on education or other issues which matter for long term development and are not able to do it in constituencies that have close elections. Our analysis would not capture this; and so in the future we hope to utilize natural experiments resulting from exogenous shocks like sudden imposition of education mandate to study the impact of educated leaders.

Even with all these limitations, our findings advance the literature on quality of leaders. We are the first, to our knowledge, to analyze the link between politician's education and competency in India. The evidence presented here suggests that having leaders with more formal education does not yield better education outcomes for citizens. Any assessment of the minimum education mandate on candidates contesting elections ought to take these findings into account. Also the literature on quality of leaders needs to consider these findings when using formal education as a proxy for quality of leaders.

5. Appendix

Table A1: OLS Estimates of the Effect on Children's Reading Score

	Reading Score					
	6–10 age-group			11–16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won by a graduate	0.027	0.010	0.034	0.061*	0.023	0.037
	(0.045)	(0.053)	(0.040)	(0.036)	(0.037)	(0.031)
Female child		0.007*	0.008*		-0.009**	-0.008*
		(0.004)	(0.004)		(0.005)	(0.005)
Child's mother went to school		0.258***	0.259***		0.221***	0.221***
		(0.006)	(0.006)		(0.006)	(0.006)
Household size		-0.037***	-0.036***		-0.030***	-0.030***
		(0.002)	(0.002)		(0.002)	(0.002)
Square of household size		0.002***	0.002***		0.001***	0.001***
		(0.000)	(0.000)		(0.000)	(0.000)
House semi-pucca		0.093***	0.090***		0.092***	0.091***
		(0.006)	(0.005)		(0.006)	(0.006)
House pucca		0.206***	0.200***		0.160***	0.158***
		(0.007)	(0.006)		(0.007)	(0.007)
Household owns television		0.141***	0.138***		0.104***	0.102***
		(0.005)	(0.005)		(0.005)	(0.004)
Household owns mobile phone		0.120***	0.129***		0.144***	0.150***
		(0.005)	(0.004)		(0.006)	(0.005)
Household's electricity use observed		0.030***	0.033***		0.034***	0.036***
		(0.005)	(0.005)		(0.005)	(0.004)
Village has electricity		0.030**	0.047***		0.027**	0.035***
		(0.012)	(0.012)		(0.012)	(0.013)
Village has pucca road		0.027***	0.028***		0.033***	0.032***
		(0.006)	(0.006)		(0.005)	(0.005)
Village has ration shop		0.025***	0.025***		0.021***	0.022***
		(0.006)	(0.006)		(0.005)	(0.005)
Village has bank		0.045***	0.042***		0.013***	0.013***
		(0.005)	(0.005)		(0.004)	(0.004)
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Age cohort fixed effects		Yes	Yes		Yes	Yes
Child, household & village controls		Yes	Yes		Yes	Yes
District controls			Yes			Yes
State by year fixed effects			Yes			Yes
Observations	1,262,927	896,820	869,544	1,272,189	928,499	904,149
Number of districts	563	563	545	563	563	545
R-squared (within)	0.005	0.051	0.058	0.002	0.040	0.044

Notes: Robust standard errors clustered at the district level are in parentheses. District controls are baseline characteristics interacted with year dummies; they include urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas estimated from 2004-05 National Sample Survey data. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table A2: OLS Estimates of the Effect on Children's Mathematics Score

| | Mathematics Score | | | | | |
| | 6–10 age-group | | | 11–16 age-group | | |
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won by a graduate	-0.017	-0.012	0.032	0.012	-0.002	0.044
	(0.049)	(0.062)	(0.041)	(0.054)	(0.066)	(0.039)
Female child		-0.043***	-0.043***		-0.084***	-0.083***
		(0.004)	(0.004)		(0.006)	(0.006)
Child's mother went to school		0.251***	0.254***		0.258***	0.259***
		(0.006)	(0.006)		(0.005)	(0.005)
Household size		-0.036***	-0.035***		-0.036***	-0.035***
		(0.002)	(0.002)		(0.002)	(0.002)
Square of household size		0.002***	0.001***		0.002***	0.002***
		(0.000)	(0.000)		(0.000)	(0.000)
House semi-pucca		0.084***	0.084***		0.095***	0.097***
		(0.005)	(0.005)		(0.006)	(0.006)
House pucca		0.211***	0.202***		0.196***	0.189***
		(0.006)	(0.006)		(0.006)	(0.006)
Household owns television		0.139***	0.138***		0.127***	0.128***
		(0.005)	(0.005)		(0.005)	(0.004)
Household owns mobile phone		0.115***	0.124***		0.126***	0.136***
		(0.004)	(0.004)		(0.005)	(0.005)
Household's electricity use observed		0.034***	0.035***		0.039***	0.041***
		(0.005)	(0.005)		(0.005)	(0.004)
Village has electricity		0.033***	0.042***		0.029**	0.018
		(0.012)	(0.012)		(0.013)	(0.013)
Village has pucca road		0.027***	0.028***		0.028***	0.029***
		(0.006)	(0.005)		(0.005)	(0.005)
Village has ration shop		0.027***	0.029***		0.020***	0.024***
		(0.006)	(0.006)		(0.005)	(0.005)
Village has bank		0.046***	0.043***		0.018***	0.016***
		(0.005)	(0.005)		(0.005)	(0.004)
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Age cohort fixed effects		Yes	Yes		Yes	Yes
Child, household & village controls		Yes	Yes		Yes	Yes
District controls			Yes			Yes
State by year fixed effects			Yes			Yes
Observations	1,251,958	892,605	865,509	1,267,391	926,214	901,972
Number of districts	563	563	545	563	563	545
R-squared (within)	0.009	0.059	0.069	0.017	0.065	0.077

Notes: Robust standard errors clustered at the district level are in parentheses. District controls are baseline characteristics interacted with year dummies; they include urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas estimated from 2004-05 National Sample Survey data. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table A3: OLS Estimates for Impact on Proportion Out of School (NSS Data)

	6-10 age-group			11-16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of	0.00371	0.00508	0.00397	-0.000318	0.00475	-0.00146
seats won by a graduate	(0.0128)	(0.0124)	(0.0113)	(0.0168)	(0.0160)	(0.0150)
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
State by year fixed effects			Yes			Yes
Observations	108521	108493	108493	131693	131651	131651
Number of districts	574	574	574	574	574	574

Robust standard errors clustered at the district level are in parentheses. District controls are interacted with year dummies and include data from 2004-05 National Sample Survey on urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table A4: First Stage Results for Impact on Proportion Out of School (NSS Data)

	6-10 age-group			11-16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of	0.963***	0.963***	0.972***	0.961***	0.961***	0.971***
seats won by a graduate in close elections	(0.111)	(0.111)	(0.112)	(0.108)	(0.107)	(0.107)
Vote margins: third order polynomial	Yes	Yes	Yes	Yes	Yes	Yes
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls		Yes	Yes		Yes	Yes
State by year fixed effects			Yes			Yes
Observations	109322	109288	109288	132815	132768	132768
Number of districts	574	574	574	574	574	574

Robust standard errors clustered at the district level are in parentheses. Close elections are defined as between a graduate and a non-graduate in which the winner beat the runner up by less than 3 percent of votes. District controls are interacted with year dummies and include data from 2004-05 National Sample Survey on urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table A5: OLS Estimates for Impact on Schools (DISE data)

	PCA for school facilites		Annual average school grants	
	(1)	(2)	(3)	(4)
Fraction of	0.0677	-0.0145	-3280.9	-4257.0**
seats won by a graduate	(0.0880)	(0.0763)	(2094.4)	(1979.9)
District fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
District controls		Yes		Yes
State by year fixed effects		Yes		Yes
Observations	2611	2381	2695	2455
Number of districts	575	521	575	521

Robust standard errors clustered at the district level are in parentheses. PCA for school facilites is generated using data on school facilites such as access to free textbooks, furniture for students, girls toilets, safe drinking water, library, playground and computers. Annual average funds refers to average total funds received by schools in the districts under School Development Grant, School Maintenance Grant, Teachers Grant and other grants. Close elections are defined as between a graduate and a non-graduate in which the winner beat the runner up by less than 3 percent of votes. District controls are interacted with year dummies and include data from 2004-05 National Sample Survey on urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table A6: First Stage Results for Impact on Schools (DISE data)

	Fraction of seats won by a graduate	
	(1)	(2)
Fraction of	0.952***	0.965***
seats won by a graduate in close elections	(0.123)	(0.130)
Vote margins: third order polynomial	Yes	Yes
District fixed effects	Yes	Yes
Year fixed effects	Yes	Yes
District controls		Yes
State by year fixed effects		Yes
Observations	2696	2242
Number of districts	575	476

Robust standard errors clustered at the district level are in parentheses. Close elections are defined as between a graduate and a non-graduate in which the winner beat the runner up by less than 3 percent of votes. District controls are interacted with year dummies and include data from 2004-05 National Sample Survey on urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table A7: 2SLS Estimates of the Effect on Children's Enrollment in Rural Areas (ASER)

	Not enrolled in school					
	6–10 age-group			11–16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won by a graduate	0.018	0.004	-0.004	0.023	0.013	0.008
	(0.012)	(0.009)	(0.009)	(0.017)	(0.016)	(0.017)
Close election fraction	0.007	0.012*	0.009	0.012	0.025**	0.013
	(0.008)	(0.006)	(0.006)	(0.011)	(0.011)	(0.011)
Female child		0.002***	0.002***		0.010***	0.010***
		(0.000)	(0.000)		(0.001)	(0.001)
Child's mother went to school		-0.013***	-0.013***		-0.050***	-0.049***
		(0.001)	(0.001)		(0.001)	(0.001)
Household size		0.002***	0.002***		0.008***	0.008***
		(0.000)	(0.000)		(0.001)	(0.001)
Square of household size		-0.000***	-0.000***		-0.000***	-0.000***
		(0.000)	(0.000)		(0.000)	(0.000)
House semi-pucca		-0.006***	-0.006***		-0.020***	-0.020***
		(0.001)	(0.001)		(0.001)	(0.001)
House pucca		-0.007***	-0.007***		-0.033***	-0.034***
		(0.001)	(0.001)		(0.002)	(0.001)
Household owns television		-0.005***	-0.005***		-0.022***	-0.021***
		(0.000)	(0.000)		(0.001)	(0.001)
Household owns mobile phone		-0.006***	-0.006***		-0.026***	-0.026***
		(0.001)	(0.001)		(0.001)	(0.001)
Household's electricity use observed		-0.002***	-0.003***		-0.008***	-0.008***
		(0.001)	(0.001)		(0.001)	(0.001)
Village has electricity		-0.012***	-0.011***		-0.010***	-0.009***
		(0.002)	(0.002)		(0.003)	(0.003)
Village has pucca road		-0.001	-0.001**		-0.006***	-0.007***
		(0.001)	(0.001)		(0.001)	(0.001)
Village has ration shop		-0.002***	-0.002***		-0.004***	-0.004***
		(0.001)	(0.001)		(0.001)	(0.001)
Village has bank		0.001***	0.001***		-0.005***	-0.005***
		(0.000)	(0.000)		(0.001)	(0.001)
Vote margins: third order polynomial	Yes	Yes	Yes	Yes	Yes	Yes
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Age cohort fixed effects		Yes	Yes		Yes	Yes
Child, household & village controls		Yes	Yes		Yes	Yes
District controls			Yes			Yes
State by year fixed effects			Yes			Yes
Observations	1,350,258	954,395	925,237	1,437,873	1,042,675	1,015,167
Number of districts	563	563	545	563	563	545
First Stage F-stat	51.33	56.62	52.67	59.68	64.08	59.28

Notes: Robust standard errors clustered at the district level are in parentheses. Close elections are defined as between a graduate and a non-graduate in which the difference in vote share between the winner and the runner up is less than 3 percent. District controls are baseline characteristics interacted with year dummies; they include urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas estimated from 2004-05 National Sample Survey data. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table A8: OLS Estimates of the Effect on Children's Grade Progression (ASER data)

	Grade Progression					
	6–10 age-group			11–16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won by a graduate	-0.024	-0.009	-0.017	-0.002	-0.007	-0.006
	(0.021)	(0.018)	(0.017)	(0.009)	(0.009)	(0.007)
Female child		0.023***	0.023***		0.005***	0.005***
		(0.002)	(0.002)		(0.001)	(0.001)
Child's mother went to school		0.013***	0.013***		0.046***	0.047***
		(0.003)	(0.003)		(0.002)	(0.002)
Household size		-0.007***	-0.007***		-0.009***	-0.009***
		(0.001)	(0.001)		(0.001)	(0.001)
Square of household size		0.000***	0.000***		0.000***	0.000***
		(0.000)	(0.000)		(0.000)	(0.000)
House semi-pucca		0.023***	0.023***		0.023***	0.022***
		(0.003)	(0.003)		(0.001)	(0.001)
House pucca		0.017***	0.018***		0.034***	0.035***
		(0.003)	(0.003)		(0.002)	(0.002)
Household owns television		0.016***	0.016***		0.023***	0.024***
		(0.003)	(0.003)		(0.001)	(0.001)
Household owns mobile phone		0.023***	0.024***		0.034***	0.034***
		(0.003)	(0.003)		(0.001)	(0.001)
Household's electricity use observed		0.011***	0.012***		0.012***	0.010***
		(0.003)	(0.003)		(0.001)	(0.001)
Village has electricity		0.055***	0.042***		0.033***	0.025***
		(0.007)	(0.007)		(0.004)	(0.004)
Village has pucca road		0.005*	0.005		0.010***	0.010***
		(0.003)	(0.003)		(0.001)	(0.001)
Village has ration shop		0.004	0.004		0.003***	0.004***
		(0.003)	(0.003)		(0.001)	(0.001)
Village has bank		-0.002	-0.003		0.004***	0.003***
		(0.003)	(0.003)		(0.001)	(0.001)
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Age cohort fixed effects		Yes	Yes		Yes	Yes
Child, household & village controls		Yes	Yes		Yes	Yes
District controls			Yes			Yes
State by year fixed effects			Yes			Yes
Observations	1,098,840	778,444	755,192	1,415,062	1,030,064	1,003,086
Number of districts	563	563	545	563	563	545
R-squared (within)	0.000	0.248	0.251	0.006	0.077	0.082

Notes: Robust standard errors clustered at the district level are in parentheses. District controls are baseline characteristics interacted with year dummies; they include urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas estimated from 2004-05 National Sample Survey data. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table A9: 2SLS Estimates of the Effect on Children's Grade Progression (ASER data)

	Grade Progression					
	6–10 age-group			11–16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won by a graduate	-0.065	0.006	0.040	-0.024	0.003	0.019
	(0.045)	(0.043)	(0.050)	(0.027)	(0.027)	(0.021)
Close election fraction	0.021	0.012	0.009	-0.015	-0.020	-0.013
	(0.024)	(0.021)	(0.024)	(0.014)	(0.015)	(0.011)
Female child		0.023***	0.023***		0.005***	0.005***
		(0.002)	(0.002)		(0.001)	(0.001)
Child's mother went to school		0.013***	0.013***		0.046***	0.047***
		(0.003)	(0.003)		(0.002)	(0.002)
Household size		-0.007***	-0.007***		-0.009***	-0.009***
		(0.001)	(0.001)		(0.001)	(0.001)
Square of household size		0.000***	0.000***		0.000***	0.000***
		(0.000)	(0.000)		(0.000)	(0.000)
House semi-pucca		0.023***	0.023***		0.023***	0.023***
		(0.003)	(0.003)		(0.001)	(0.001)
House pucca		0.017***	0.017***		0.035***	0.035***
		(0.003)	(0.003)		(0.002)	(0.002)
Household owns television		0.015***	0.016***		0.023***	0.024***
		(0.003)	(0.003)		(0.001)	(0.001)
Household owns mobile phone		0.023***	0.024***		0.034***	0.034***
		(0.003)	(0.003)		(0.001)	(0.001)
Household's electricity use observed		0.011***	0.011***		0.012***	0.010***
		(0.003)	(0.002)		(0.001)	(0.001)
Village has electricity		0.054***	0.042***		0.032***	0.025***
		(0.007)	(0.007)		(0.004)	(0.004)
Village has pucca road		0.005*	0.005*		0.010***	0.010***
		(0.003)	(0.003)		(0.001)	(0.001)
Village has ration shop		0.004	0.005		0.003***	0.004***
		(0.003)	(0.003)		(0.001)	(0.001)
Village has bank		-0.002	-0.003		0.004***	0.003***
		(0.003)	(0.003)		(0.001)	(0.001)
Vote margins: third order polynomial	Yes	Yes	Yes	Yes	Yes	Yes
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Age cohort fixed effects		Yes	Yes		Yes	Yes
Child, household & village controls		Yes	Yes		Yes	Yes
District controls			Yes			Yes
State by year fixed effects			Yes			Yes
Observations	1,098,840	778,444	755,192	1,415,062	1,030,064	1,003,086
Number of districts	563	563	545	563	563	545
First Stage F-stat	51.27	56.38	52.32	59.84	64.24	59.34

Notes: Robust standard errors clustered at the district level are in parentheses. Close elections are defined as between a graduate and a non-graduate in which the difference in vote share between the winner and the runner up is less than 3 percent. District controls are baseline characteristics interacted with year dummies; they include urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas estimated from 2004-05 National Sample Survey data.
*** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table A10: OLS Estimates of the Effect on Children's English Score

	English Score					
	6–10 age-group			11–16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won by a graduate	0.036	0.067	0.087	0.024	-0.002	0.025
	(0.056)	(0.075)	(0.067)	(0.054)	(0.067)	(0.054)
Female child		-0.043***	-0.043***		-0.062***	-0.062***
		(0.005)	(0.005)		(0.006)	(0.006)
Child's mother went to school		0.313***	0.315***		0.290***	0.291***
		(0.007)	(0.007)		(0.007)	(0.007)
Household size		-0.034***	-0.035***		-0.041***	-0.041***
		(0.003)	(0.003)		(0.003)	(0.003)
Square of household size		0.002***	0.002***		0.002***	0.002***
		(0.000)	(0.000)		(0.000)	(0.000)
House semi-pucca		0.081***	0.079***		0.092***	0.090***
		(0.006)	(0.006)		(0.007)	(0.007)
House pucca		0.265***	0.260***		0.226***	0.221***
		(0.008)	(0.008)		(0.009)	(0.008)
Household owns television		0.173***	0.174***		0.146***	0.144***
		(0.006)	(0.006)		(0.006)	(0.006)
Household owns mobile phone		0.142***	0.149***		0.158***	0.167***
		(0.006)	(0.006)		(0.007)	(0.007)
Household's electricity use observed		0.026***	0.026***		0.035***	0.031***
		(0.007)	(0.007)		(0.007)	(0.006)
Village has electricity		-0.019	0.010		-0.005	0.002
		(0.018)	(0.018)		(0.022)	(0.022)
Village has pucca road		0.033***	0.040***		0.037***	0.044***
		(0.009)	(0.009)		(0.008)	(0.008)
Village has ration shop		0.043***	0.040***		0.029***	0.032***
		(0.008)	(0.008)		(0.008)	(0.007)
Village has bank		0.078***	0.076***		0.029***	0.027***
		(0.007)	(0.008)		(0.007)	(0.007)
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Age cohort fixed effects		Yes	Yes		Yes	Yes
Child, household & village controls		Yes	Yes		Yes	Yes
District controls			Yes			Yes
State by year fixed effects			Yes			Yes
Observations	554,483	367,406	355,908	563,175	388,105	378,254
Number of districts	563	561	544	563	561	544
R-squared (within)	0.001	0.086	0.093	0.002	0.070	0.078

Notes: Robust standard errors clustered at the district level are in parentheses. District controls are baseline characteristics interacted with year dummies; they include urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas estimated from 2004-05 National Sample Survey data. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table A11: 2SLS Estimates of the Effect on Children's English Score

	English Score					
	6–10 age-group			11–16 age-group		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of seats won by a graduate	0.040	0.266	0.187	-0.062	0.119	0.175
	(0.152)	(0.176)	(0.157)	(0.166)	(0.173)	(0.136)
Close election fraction	0.099	0.065	0.016	0.058	0.024	-0.024
	(0.079)	(0.081)	(0.071)	(0.084)	(0.083)	(0.079)
Female child		-0.042***	-0.043***		-0.062***	-0.061***
		(0.004)	(0.005)		(0.006)	(0.006)
Child's mother went to school		0.312***	0.313***		0.290***	0.290***
		(0.007)	(0.007)		(0.007)	(0.007)
Household size		-0.034***	-0.035***		-0.041***	-0.041***
		(0.003)	(0.003)		(0.003)	(0.003)
Square of household size		0.002***	0.002***		0.002***	0.002***
		(0.000)	(0.000)		(0.000)	(0.000)
House semi-pucca		0.082***	0.081***		0.091***	0.090***
		(0.006)	(0.006)		(0.007)	(0.007)
House pucca		0.266***	0.261***		0.226***	0.221***
		(0.008)	(0.008)		(0.009)	(0.008)
Household owns television		0.173***	0.173***		0.146***	0.144***
		(0.006)	(0.006)		(0.006)	(0.006)
Household owns mobile phone		0.141***	0.148***		0.158***	0.166***
		(0.006)	(0.006)		(0.007)	(0.007)
Household's electricity use observed		0.027***	0.026***		0.035***	0.030***
		(0.007)	(0.007)		(0.007)	(0.006)
Village has electricity		-0.017	0.007		-0.004	-0.001
		(0.018)	(0.018)		(0.022)	(0.023)
Village has pucca road		0.034***	0.042***		0.038***	0.046***
		(0.009)	(0.009)		(0.008)	(0.008)
Village has ration shop		0.042***	0.040***		0.030***	0.033***
		(0.008)	(0.008)		(0.007)	(0.007)
Village has bank		0.077***	0.076***		0.030***	0.027***
		(0.008)	(0.008)		(0.007)	(0.007)
Vote margins: third order polynomial	Yes	Yes	Yes	Yes	Yes	Yes
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Age cohort fixed effects		Yes	Yes		Yes	Yes
Child, household & village controls		Yes	Yes		Yes	Yes
District controls			Yes			Yes
State by year fixed effects			Yes			Yes
Observations	554,483	367,406	355,908	563,175	388,105	378,254
Number of districts	563	561	544	563	561	544
First Stage F-stat	64.76	69.21	63.71	70.49	72.68	66.67

Notes: Robust standard errors clustered at the district level are in parentheses. Close elections are defined as between a graduate and a non-graduate in which the difference in vote share between the winner and the runner up is less than 3 percent. District controls are baseline characteristics interacted with year dummies; they include urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas estimated from 2004-05 National Sample Survey data. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table A12: Descriptive Statistics for Schooling Outcomes (ASER data)

	Mean	Std.Dev.	Obs
Usable blackboard in class 2	0.94	0.23	40526
Usable blackboard in class 4	0.94	0.24	40526
Learning material in class 2	0.80	0.40	40526
Learning material in class 4	0.77	0.42	40526
Drinking water	0.76	0.42	40526
Usable toilet	0.79	0.41	40526
Usable library books	0.44	0.50	40526
PCA score for physical assets	0.04	1.50	40526
Schoool maintenance grant	0.89	0.31	50184
School development grant	0.82	0.39	50184
Teacher learning material grant	0.72	0.45	50184
PCA score SSA school grants	–0.03	1.28	50184
Mid-day meal	0.85	0.36	77692
Children's attendance	0.72	0.20	76730
Teacher's attendance	0.85	0.25	61641

Table A13: Impact on School Attendance (ASER data)

	Child Attendance Ratio		Teacher Attendance Ratio	
	(1)	(2)	(3)	(4)
Fraction of	-0.0216	-0.0170	0.0527	0.0751
seats won by a graduate	(0.0303)	(0.0338)	(0.0487)	(0.0519)
Vote margins: third order polynomial	Yes	Yes	Yes	Yes
District fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
School type controls		Yes		Yes
Village controls		Yes		Yes
District controls		Yes		Yes
State by year fixed effects		Yes		Yes
Observations	76730	69906	61641	56096
Number of districts	563	544	563	544
First stage F statistics	67	61	71	62

Robust standard errors clustered at the district level are in parentheses. Close elections are defined as between a graduate and a non-graduate in which the winner beat the runner up by less than 3 percent of votes. Controls at school level include dummies for if school is upto primary only - upto 4-5 or includes middle school - upto 7-8 or other. Controls for village include dummies for existence of electricity, permanent road, ration shop and bank. District controls are interacted with year dummies and include data from 2004-05 National Sample Survey on urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

Table A14: Impact on Schooling Inputs (ASER data)

	PCA for physical assets		PCA for grants		Midday meals	
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction of	-0.207	-0.0859	-0.550*	-0.386	0.0423	0.103
seats won by a graduate	(0.303)	(0.317)	(0.282)	(0.240)	(0.0739)	(0.0738)
Vote margins: third order polynomial	Yes	Yes	Yes	Yes	Yes	Yes
District fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
School type controls		Yes		Yes		Yes
Village controls		Yes		Yes		Yes
District controls		Yes		Yes		Yes
State by year fixed effects		Yes		Yes		Yes
Observations	56262	51273	50184	47149	77692	69739
Number of districts	563	542	561	544	563	544
First stage F statistics	73	60	74	70	68	60

Robust standard errors clustered at the district level are in parentheses. Close elections are defined as between a graduate and a non-graduate in which the winner beat the runner up by less than 3 percent of votes. Controls at school level include dummies for if school is upto primary only - upto 4-5 or includes middle school - upto 7-8 or other. Controls for village include dummies for existence of electricity, permanent road, ration shop and bank. District controls are interacted with year dummies and include data from 2004-05 National Sample Survey on urbanization rate, proportion of adult male and adult female in rural areas who have completed primary school, caste composition and proportion of females in rural areas. *** Significant at the 1 percent level. ** Significant at the 5 percent level. * Significant at the 10 percent level.

References

Aaberg, R. and I. Melby (1998), The Sensitivity of Income Inequality to Choice of Equivalence Scales, *Review of Income and Wealth*, Vol. 44, No. 4.

Abdul, T. (2015). Dismantling democracy: How the supreme court deprived the poor of their right to fight elections. http://www.huffingtonpost.in/talha-abdul-rahman-/dismantlingdemocracy-how_b_9137248.html?m=false.

Alderman, H., Chiappori, P. A., Haddad, L., Hoddinott, J., & Kanbur, R. (1995). Unitary versus collective models of the household: is it time to shift the burden of proof? *The World Bank Research Observer*, *10*(1), 1-19.

Alkire, S., & Foster, J. (2011). Counting and multidimensional poverty measurement. *Journal of public economics*, *95*(7), 476-487.

Alkire, S., Meinzen-Dick, R., Peterman, A., Quisumbing, A., Seymour, G., & Vaz, A. (2013). The women's empowerment in agriculture index. *World Development*, *52*, 71-91.

Alkire, S., & Santos, M. E. (2014). Measuring acute poverty in the developing world: Robustness and scope of the multidimensional poverty index. *World Development*, *59*, 251-274.

Alkire, S., & Seth, S. (2015). Multidimensional Poverty Reduction in India between 1999 and 2006: Where and How?. *World Development*, *72*, 93-108.

Anand, S. and P. Morduch (1996), Poverty and the Population Problem: Evidence from Bangladesh, Harvard Institute for International Development, Development Discussion Paper No. 559, available on http://www.cid.harvard.edu/hiid/559.pdf .

Anand, S., & P. Segal. (2008). What do we Know about Global Income Inequality? *Journal of Economic Literature*, 57-94.

Anand, S., and Segal, P. (2014). The Global Distribution of Income. Handbook of Income Distribution. Volume 2A. Ch. 11, pp. 937-979.

Aronsson, T., Daunfeldt, S. O., & Wikström, M. (2001). Estimating intrahousehold allocation in a collective model with household production. *Journal of Population Economics*, *14*(4), 569-584.

Asfaw, A., Lamanna, F., & Klasen, S. (2010). Gender gap in parents' financing strategy for hospitalization of their children: evidence from India. *Health economics*, *19*(3), 265-279.

Asher, S. and P. Novosad (2013). Politics and local economic growth: Evidence from india. American Economic Journal: Applied Economics.

Atkinson, A, & A. Brandolini. (2001). Promise and Pitfalls in the Use of "Secondary" Data-Sets: Income Inequality in OECD Countries As a Case Study. *Journal of Economic Literature*, Vol. 39, No.3, pages 771-799, September.

Atkinson, A. B., & A. Brandolini, (2009). On data: a case study of the evolution of income inequality across time and across countries. *Cambridge Journal of Economics*, Vol. *33* No. 3, 381-404.

Atkinson, A. B., & Brandolini, A. (2010). On analyzing the world distribution of income. *The World Bank Economic Review*, lhp020.

Atkinson, M. M., D. Rogers, and S. Olfert (2016). Better politicians: If we pay, will they come? Legislative Studies Quarterly 41(2), 361–391.

193

Banerjee, A., S. Cole, E. Duflo, and L. Linden (2005). Remedying education: Evidence from two randomized experiments in india. Technical report, National Bureau of Economic Research.
Banerjee, A. V., R. Banerji, E. Duflo, R. Glennerster, and S. Khemani (2010). Pitfalls of participatory programs: Evidence from a randomized evaluation in education in india. American Economic Journal: Economic Policy 2(1), 1–30.

Banerji, R., J. Berry, and M. Shotland (2013). The impact of mother literacy and participation programs on child learning: Evidence from a randomized evaluation in india. Cambridge, MA: Abdul Latif Jameel Poverty Action Lab (J-PAL).

Banks, J. and P. Johnson (1994). Equivalence Scale Relativities Revisited, *Economic Journal*, Vol. 104, pp. 883-890.

Basu, K., & Foster, J. E. (1998). On measuring literacy. *The Economic Journal*, *108*(451), 1733-1749.

Baxi, U. (2015). Supreme error. http://indianexpress.com/article/opinion/columns/supremeerror/.

Besley, T. and S. Coate (1997). An economic model of representative democracy. The Quarterly Journal of Economics, 85–114.

Besley, T., J. G. Montalvo, and M. Reynal-Querol (2011). Do educated leaders matter? The Economic Journal 121(554), F205–227.

Besley, T., R. Pande, and V. Rao (2005). Political selection and the quality of government: Evidence from south india. Yale University Economic Growth Center Discussion Paper No. 921.

Besley, T. and M. Reynal-Querol (2011). Do democracies select more educated leaders? American political science review 105(03), 552–566.

Bessell, S. (2015). The Individual Deprivation Measure: measuring poverty as if gender and inequality matter. *Gender & Development*, *23*(2), 223-240.

Bhalla, S. (2002). Imagine there's no country: Poverty, inequality, and growth in the era of globalization. Peterson Institute.

Bhalotra, S. and I. Clots-Figueras (2014). Health and the political agency of women. American Economic Journal: Economic Policy 6(2), 164–197.

Bhalotra, S., I. Clots-Figueras, G. Cassan, and L. Iyer (2014). Religion, politician identity and development outcomes: Evidence from india. Journal of Economic Behavior & Organization 104, 4–17.

Bhalotra, S. R., I. Clots-Figueras, and L. Iyer (2013). Path-breakers: How does women's political participation respond to electoral success?

Bhaskar, A. (2016). Damage to democracy. Economic & Political Weekly 51(40), 47.

Bhavnani, R. R. and F. Jensenius (2015). Voting for development? ruling coalitions and literacy in india. Technical report, Working paper. Madison and Oslo: University of Wisconsin–Madison and Norwegian Institute of International Affairs.

Blaylock (1991). The Impact of Equivalence Scales on the Analysis of Income and Food Spending Distributions. *Western Journal of Agricultural Economics*, Vol. 16, No 1, pp. 11-20.

Bossert, W., Chakravarty, S. R., & D'Ambroiso, C. (2009). Multidimensional poverty and material deprivation.

Bourguignon, F. (2012): La mondialisation de l'inégalité. Seuil et La République des Idées.

194

Bourguignon, F., and Morrisson, C. (2002): Inequality Among World Citizens: 1820-1992. American Economic Review, Vol. 92(4), pp.727-744.

Browning, M., Chiappori, P. A., & Lewbel, A. (2013). Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *The Review of Economic Studies*, rdt019.

Buhmann, B, L. Rainwater, G. Schmaus and T. Smeeding (1988). Equivalence Scales, Well-Being, Inequality, And Poverty: Sensitivity Estimates Across Ten Countries Using The Luxembourg Income Study (Lis) Database. *Review of Income and Wealth*, Vol.34, No.2, pp. 115-142.

Burchi, F. (2013). Women's political role and poverty in the educational dimension. a districtlevel analysis in india. German Development Institute Discussion Paper (23).

Case, A., & Deaton, A. (2003). *Consumption, health, gender, and poverty* (Vol. 3020). World Bank Publications.

Carnes, N. (2012). Does the numerical underrepresentation of the working class in congress matter? Legislative Studies Quarterly 37(1), 5–34.

Carnes, N. (2013). White-collar government: The hidden role of class in economic policy making. University of Chicago Press.

Carnes, N. and N. Lupu (2016). What good is a college degree? education and leader quality reconsidered. The Journal of Politics 78(1), 35–49.

Caughey, D. and J. S. Sekhon (2011). Elections and the regression discontinuity design: Lessons from close us house races, 1942-2008. Political Analysis, 385–408.

Chakravarty, S. R., & D'Ambrosio, C. (2006). The measurement of social exclusion. *Review of income and wealth*, *52*(3), 377-398.

Chant, S. (2004). Dangerous Equations? How Female- headed Households Became the Poorest of the Poor: Causes, Consequences and Cautions. *IDS bulletin*, *35*(4), 19-26.

Chattopadhyay, R. and E. Duflo (2004). Women as policy makers: Evidence from a randomized policy experiment in india. Econometrica 72(5), 1409–1443.

Chen, S., and M. Ravallion (2004). How Have the World's Poorest Fared Since the Early 1980s?. *World Bank Research Observer*, Vol. 19, No. 2, pp. 141-169.

Chiappori, P. A. (1988). Rational household labor supply. *Econometrica: Journal of the Econometric Society*, 63-90.

Chiappori, P. A. (1992). Collective labor supply and welfare. *Journal of political Economy*, 437-467.

Chiappori, P. A., Fortin, B., & Lacroix, G. (2002). Marriage market, divorce legislation, and household labor supply. *Journal of political Economy*, *110*(1), 37-72.

Chiappori, P. A., & Meghir, C. (2014). *Intrahousehold inequality* (No. w20191). National Bureau of Economic Research.

Chotikapanich, D. (1993). A comparison of Alternative Functional Forms for the Lorenz Curve. *Economics Letters*, Vol. *41* No. 2, 129-138.

Chopra, V. K. (1996). Marginal players in marginal assemblies: The Indian MLA. Orient Longman.

Chotikapanich, D., Valenzuela, R., and Prasada Rao, D. S. (1997). Global and Regional Inequality in the Distribution of Income: Estimation With Limited/Incomplete Data. Empirical Economics, Vol. 22(4), pp. 533–546.

Clotfelter, C. T., H. F. Ladd, and J. L. Vigdor (2007). Teacher credentials and student achievement: Longitudinal

analysis with student fixed effects. Economics of Education Review 26(6), 673–682.

Clots-Figueras, I. (2011). Women in politics: Evidence from the indian states. Journal of public Economics 95(7), 664–690.

Clots-Figueras, I. (2012). Are female leaders good for education? American Economic Journal: Applied Economics 4(1), 212–244.

Congleton, R. D. and Y. Zhang (2013). Is it all about competence? the human capital of us presidents and economic performance. Constitutional Political Economy 24(2), 108–124.

Corak, M., Fertig, M., & Tamm, M. (2008). A portrait of child poverty in Germany. *Review of Income and Wealth*, *54*(4), 547-571.

Corlett, Adam (2016)**.** Examining an elephant: globalisation and the lower middle class of the rich world. Resolution Foundation. Available at : http://www.resolutionfoundation.org/publications/examining-an-elephant-globalisation-and-the-lower-middle-class-of-the-rich-world/

Coulter, F. A. E., Cowell, F. A. and Jenkins, S. P. (1992a). Differences in Needs and Assessment of Income Distributions. *Bulletin of Economic Research*, Vol. 44, pp. 77-124

Coulter, F. A. E., Cowell, F. A. and Jenkins, S. P. (1992b). Equivalence Scale Relativities and the Extent of Inequality and Poverty. *Economic Journal*. Vol. 102, pp.1067-1082.

Cowell, F.A. and M. Mercader-Prats (1999). Equivalence Scales and Inequality. LSE STICERD Working Paper No. 27, available on http://eprints.lse.ac.uk/2190/1/Equivalence_of_Scales_and_Inequality.pdf

Datt, G. (1998). Computational Tools for Poverty Measurement and Analysis. Washington, D.C.: International Food Policy Research Institute.

de Canberra, G. (2001). Expert group on household Income statistics: Final report and recommendations.

Deaton, A. (2005). Measuring Poverty in a Growing World (or Measuring Growth in a Poor World). *Review of Economics and Statistics*. *87*(2), 395-395.

Deaton, A., & Aten, B. (2017). Trying to Understand the PPPs in ICP 2011: Why Are the Results So Different? *American Economic Journal: Macroeconomics*, *9*(1), 243-64.

Deaton, A., and A. Heston. (2010). Understanding PPPs and PPP-based National Accounts. American Economic Journal: Macroeconomics 2 (4): 1–35.

Deaton, A. S., & Paxson, C. (1998). Measuring Poverty among the Elderly. In *Inquiries in the Economics of Aging* (pp. 169-204). University of Chicago Press.

Decancq, K., Van Ootegem, L., & Verhofstadt, E. (2013). What If We Voted on the Weights of a Multidimensional Well- Being Index? An Illustration with Flemish Data. *Fiscal Studies*, *34*(3), 315-332.

Deininger, K., and L. Squire (1996). A New Data Set Measuring Income Inequality. *The World Bank Economic Review*, *10*(3), 565-591.

Desai, S. B., Dubey, A., Joshi, B. L., Sen, M., Shariff, A., & Vanneman, R. (2010). *Human development in India*. Oxford University, New York.

Dhongde, S. and M. Xing. (2013). Cross-Country Convergence in Income Inequality. Working Papers 290, ECINEQ, Society for the Study of Economic Inequality.

Diaz-Serrano, L. and J. Pérez (2013). Do more educated leaders raise citizens' education?

Dikhanov, Y., and Ward, M. (2002): Evolution of the Global Distribution of Income 1970–99. 53rd Session of the International Statistical Institute, Seoul, Republic of Korea, pp 22- 29.

Dotter, C., & Klasen, S. (2014). The Multidimensional Poverty Index: Achievements, Conceptual, and Empirical Issues. *Human Development Report Office Occasional Paper*.

Downs, A. (1957). An economic theory of political action in a democracy. The journal of political economy, 135–150.

Dowrick, S., and Akmal, M. (2005). Contradictory Trends in Global Income Inequality: A Tale of Two Biases. Review of Income and Wealth. Vol. 51(2): pp. 201-229.

Dreher, A., M. J. Lamla, S. M. Lein, and F. Somogyi (2009). The impact of political leaders' profession and education on reforms. Journal of Comparative Economics 37(1), 169–193.

Dreze, J., & Srinivasan, P. V. (1997). Widowhood and poverty in rural India: Some inferences from household survey data. *Journal of Development Economics*, *54*(2), 217-234.

Duclos, J. Y. (2011). UNDP's Multidimensional poverty index. *Development*, 11.

Dunbar, G. R., Lewbel, A., & Pendakur, K. (2013). Children's resources in collective households: identification, estimation, and an application to child poverty in Malawi. *The American Economic Review*, *103*(1), 438-471.

Dykstra, S., B. Dykstra, B. and J. Sandefur, (2014). We Just Ran 23 Million Queries of the World Bank's Website. Center for Global Development, Working Paper 362.

Edlund, L., & Korn, E. (2002). A theory of prostitution. *Journal of political Economy*, *110*(1), 181-214.

Edward, P and Sumner A (2013). The Geography of Inequality: Where and by How Much Has Income Distribution Changed since 1990? Center for Global Development, Working Paper 362

Edward, P. and Sumner A. (2015). *New Estimates of Global Poverty and Inequality: How Much Difference Do Price Data Really Make? Working Paper 403* (No. 403). Center for Global Development.

Eggers, A. C., A. Fowler, J. Hainmueller, A. B. Hall, and J. M. Snyder (2015). On the validity of the regression discontinuity design for estimating electoral effects: New evidence from over 40,000 close races. American Journal of Political Science 59(1), 259–274.

Ezzati, M., & Kammen, D. M. (2002). The health impacts of exposure to indoor air pollution from solid fuels in developing countries: knowledge, gaps, and data needs. *Environmental health perspectives*, *110*(11), 1057.

Ferreira, F. H., Chen, S., Dabalen, A., Dikhanov, Y., Hamadeh, N., Jolliffe, D., ... & Serajuddin, U. (2016). A global count of the extreme poor in 2012: data issues, methodology and initial results. *The Journal of Economic Inequality*, *14*(2), 141-172.

Gaiha, R., Jha, R., & Kulkarni, V. S. (2010). Child Undernutrition in India. *Available at SSRN 1734591*.

Gersbach, H., & Haller, H. (2001). Collective decisions and competitive markets. *The Review of Economic Studies*, *68*(2), 347-368.

Ghani, S. E., A. Mani, and S. D. O'Connell (2013). Can political empowerment help economic empowerment? women leaders and female labor force participation in india. World Bank Policy Research Working Paper (6675).

Gottesman, A. A. and M. R. Morey (2006). Manager education and mutual fund performance. Journal of empirical finance 13(2), 145–182.

Grimmer, J., E. Hersh, B. Feinstein, and D. Carpenter (2011). Are close elections random? Unpublished manuscript.

Grossbard-Shechtman, S. (1993). On the economics of marriage: a theory of marriage labor and divorce.

Haddad, L., & Kanbur, R. (1990). How Serious Is the Neglect of Intra-Household Inequality? *Economic Journal*, *100*(402), 866-81.

Haddad, L., Hoddinott, J., & Alderman, H. (1997). *Intrahousehold resource allocation in developing countries: models, methods, and policy*. Johns Hopkins University Press.

Halim, N., K. M. Yount, S. A. Cunningham, and R. P. Pande (2016). Women's political empowerment and investments in primary schooling in india. Social indicators research 125(3), 813–851.

Hazarika, G. (2000). Gender differences in children's nutrition and access to health care in Pakistan. *The Journal of Development Studies*, *37*(1), 73-92.

Imbens, G. W. and T. Lemieux (2008). Regression discontinuity designs: A guide to practice. Journal of econometrics 142(2), 615–635.

International Institute for Population Sciences (IIPS) and Macro International. 2007. National Family Health Survey (NFHS-3), 2005–06: India: Volume I.

Iyer, L., A. Mani, P. Mishra, and P. Topalova (2012). The power of political voice: women's political representation and crime in india. American Economic Journal: Applied Economics 4(4), 165–193.

Jaffrelot, C. (2016). Laws prescribing educational qualifications for contesting elections are undemocratic. http://indianexpress.com/article/opinon/columns/narendra-modi-pmdegree-arvind-kejriwal-delhi-university-ba-degrees-of-exclusion-2792374/.

Jain, S. (1975). *Size distribution of income: A compilation of data* (pp. 101-104). Washington, DC: World Bank.

Jayadev, A., Lahoti, R., and Reddy, S. G. (2015a). Who Got What, Then and Now? A Fifty Years Overview from the Global Consumption and Income Project.: working paper, May 4th.

Jayadev, A., Lahoti, R., and Reddy, S. G. (2015b). The Middle Muddle: Conceptualizing and Measuring the Global Middle Class. working paper.

Jayaraj, D., & Subramanian, S. (2010). A Chakravarty-D'Ambrosio view of multidimensional deprivation: some estimates for India. *Economic and Political Weekly*, 53-65.

Jensenius, F. R. (2015). Development from representation? a study of quotas for the scheduled castes in india. American Economic Journal: Applied Economics 7(3), 196–220.

Jensenius, F. R. and P. Suryanarayan (2015). Fragmentation and decline in india's state assemblies. Asian Survey 55(5), 862–881.

Jenkins, S. P. (2015). World Income Inequality Databases: an assessment of WIID and SWIID, *The Journal of Economic Inequality*

Jenkins, S. P., & Lambert, P. J. (1997). Three 'I's of poverty curves, with an analysis of UK poverty trends. *Oxford Economic Papers*, *49*(3), 317-327.

Jones, B. F. and B. A. Olken (2009). Hit or miss? the effect of assassinations on institutions and war. American Economic Journal: Macroeconomics, 55–87.

Kakwani, N.C. (1980). Functional Forms for Estimating the Lorenz Curve: A Reply. *Econometrica*, Vol. 48, Issue 4, pp 1063-1064

Klasen, S. (2007). *Gender-related indicators of well-being* (pp. 167-192). Palgrave Macmillan UK.

Klasen, S. (2008). Poverty, undernutrition, and child mortality: Some inter-regional puzzles and their implications for research and policy. *The Journal of Economic Inequality*, *6*(1), 89-115.

Klasen, S., Lechtenfeld, T., & Povel, F. (2015). A feminization of vulnerability? Female headship, poverty, and vulnerability in Thailand and Vietnam. *World Development*, *71*, 36-53.

Klasen, S., Krivobokova, T., Greb, F., Lahoti, R., Pasaribu, S. H., & Wiesenfarth, M. (2016). International income poverty measurement: which way now? *The Journal of Economic Inequality*, *14*(2), 199-225.

Klasen, S., Scholl, N., Lahoti, R., Ochmann, S., & Vollmer, S. (2016). *Inequality-worldwide trends and current debates* (No. 209). Courant Research Centre: Poverty, Equity and Growth-Discussion Papers.

Klasen, S., & Wink, C. (2002). A turning point in gender bias in mortality? An update on the number of missing women. *Population and Development Review*, *28*(2), 285-312.

Klasen, S., & Wink, C. (2003). "Missing women": Revisiting the debate. *Feminist Economics*, *9*(2-3), 263-299.

Krishnan, K. and T. Somanathan (2013). The civil service in the 21st century. Unpublished manuscript. New Delhi, India.

Lahoti, R., A. Jayadev and S.G. Reddy (2014). The Global Consumption and Income Project (GCIP): An Introduction and Preliminary Findings. SSRN working paper available on http://ssrn.com/abstract=2480636

Lahoti, R., Suchitra, J. Y., & Swaminathan, H. (2016). Not in Her Name. *Economic & Political Weekly*, *51*(5), 17.

Lakner, C., & Milanovic, B. (2016). Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession. *World Bank Economic Review*, *30*(2), 203-232.

Lakner, C. and B. Milanovic (2016). Response to Adam Corlett's Examining an elephant: globalisation and the lower middle class of the rich world. Mimeo CUNY

Lasswell, H. (1936), *Politics: Who Gets What, When. How*. New York: Whittlesey. House. 1936.

Levitt, S. D. (1996). How do senators vote? disentangling the role of voter preferences, party affliation, and senator ideology. The American Economic Review, 425–441.

Li, H., L. Squire, and H.-F. Zou. (1998). Explaining International and Intertemporal Variations in Income Inequality. The Economic Journal 108 (446): 26–43.

Lise, J., & Seitz, S. (2011). Consumption inequality and intra-household allocations. *The Review of Economic Studies*, *78*(1), 328-355.

Lundberg, S., & Pollak, R. A. (1993). Separate spheres bargaining and the marriage market. *Journal of political Economy*, 988-1010

Luxembourg Income Study Database (LIS), www.lisdatacenter.org (multiple countries; June 1st- June 3rd 2014). Luxembourg: LIS.

Manser, M., & Brown, M. (1980). Marriage and household decision-making: A bargaining analysis. *International economic review*, 31-44.

Martinez-Bravo, M. (2014). The local political economy effects of school construction in indonesia. American Economic Journal: Applied Economics.

Marcoux, A. (1998). The feminization of poverty: claims, facts, and data needs. *Population and development review*, 131-139.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. Journal of Econometrics 142(2), 698–714.

McElroy, M. B., & Horney, M. J. (1981). Nash-bargained household decisions: Toward a generalization of the theory of demand. *International economic review*, 333-349

Milanovic, B. (2002). True World Income Distribution, 1988 And 1993: First Calculation Based On Household Surveys Alone. *The Economic Journal*, *112*(476), 51-92.

Milanovic, B. (2005). *Worlds Apart: Measuring International and Global Inequality*. Princeton, N.J.: Princeton University Press.

Milanovic, B. (2012). Global Income Inequality by the Numbers in History and Now - an Overview. Washington, D.C.: World Bank, Development Research Group, Poverty and Inequality Team.

Minoiu, C. and S.G. Reddy (2009). The Estimation of Poverty and Inequality through Parametric Estimation of Lorenz Curves: An Evaluation. *Journal of Income Distribution*, vol. 18(2), pages 160-178, June.

Nath, A. (2015). Bureaucrats and politicians: How does electoral competition affect bureaucratic performance? Unpublished manuscript, Boston University.

NDTV (2016). Does education make for better politicians? http://www.ndtv.com/video/news/we-the-people/does-education-make-for-better-politicians-411370.

Niño- Zarazúa, M., Roope, L., & Tarp, F. (2016). Global inequality: Relatively lower, absolutely higher. *Review of Income and Wealth*.

Niño Zarazúa, M., Roope, L., & Tarp, F. (2014). Global Interpersonal Inequality: Trends and Measurement

Nurkse, R. (1953). The Problem of capital Formation in less-developed Countries. *Oxford University Press*, *33*, 1-337.

Nussbaum, M. (2003). Capabilities as fundamental entitlements: Sen and social justice. *Feminist economics*, *9*(2-3), 33-59.

O'Connell, S. D. (2015). Political inclusion and educational investment.

Osborne, M. J. and A. Slivinski (1996). A model of political competition with citizencandidates. The Quarterly Journal of Economics, 65–96.

Pinkovskiy, M., & Sala-i-Martin, X. (2009). Parametric Estimations of the World Distribution of Income. NBER Working Paper 15433.

Prakash, N., M. Rockmore, and Y. Uppal (2014). Do criminally accused politicians affect economic outcomes? evidence from india. Technical report, Households in Conflict Network.

Pritchett, L. (2013). The rebirth of education: Schooling ain't learning. CGD Books.

Quah, D.T. (1996). Twin Peaks: Growth and Convergence in Models of Distribution Dynamics. *Economic Journal*, vol. 106(437), pages 1045-55, July.

Quisumbing, A. R., & Maluccio, J. A. (2000). *Intrahousehold allocation and gender relations: New empirical evidence from four developing countries*. Washington, DC: International Food Policy Research Institute.

Rasche, R. H, Gaffney, J, A.Y.C Koo, and N Obst (1980). Functional Forms for Estimating the Lorenz Curve, *Econometrica* Vol. 48, No. 4 (May, 1980) , pp. 1061-1062

Ravallion, M. (2003). Inequality Convergence. *Economics Letters*, Elsevier, vol. 80(3), pages 351-356, September.

Ravallion, M. (2003b). Measuring aggregate welfare in developing countries: How well do national accounts and surveys agree?. *Review of Economics and Statistics*, *85*(3), 645-652.

———. 2014. An Exploration of the International Comparison Program's New Global Economic Landscape. Working paper, Georgetown University

Reddy, S. G., and Lahoti, R. (2015). $1.90 per day: What does it say? (No. 189). Courant Research Centre: Poverty, Equity and Growth-Discussion Papers.

Reddy, S. and T. Pogge (2010). How Not to Count the Poor, in Segal, P., S. Anand and J. E. Stiglitz ed., *Debates on the Measurement of Global Poverty*, Oxford: Oxford University Press.

Rippin, N. (2012). *Distributional justice and efficiency: Integrating inequality within and between dimensions in additive poverty indices* (No. 128). Courant Research Centre: Poverty, Equity and Growth-Discussion Papers.

Roche, J. M. (2013). Monitoring progress in child poverty reduction: Methodological insights and illustration to the case study of Bangladesh. *Social indicators research*, *112*(2), 363-390.

Rosenzweig, M. R., & Schultz, T. P. (1982). Market opportunities, genetic endowments, and intrafamily resource distribution: Child survival in rural India. *The American Economic Review*, *72*(4), 803-815.

Sala-i-Martin, X. (2006): The World Distribution of Income: Falling Poverty and... Convergence, Period. The Quarterly Journal of Economics. Vol. 121(2): pp. 351-397.

Sefil, S. (2015). Sensitivity Of Turkish Income Distributions to Choice of Equivalence Scale. *Topics in Middle Eastern and African Economies*, Vol. 17, Issue No. 1.

Sen, A. (2002). Health: perception versus observation : Self reported morbidity has severe limitations and can be extremely misleading. *BMJ : British Medical Journal*, *324*(7342), 860–861.

Sen, A. (1976). Poverty: an ordinal approach to measurement. *Econometrica: Journal of the Econometric Society*, 219-231.

Serajuddin, Umar; Uematsu, Hiroki; Wieser, Christina; Yoshida, Nobuo; Dabalen, Andrew L.. 2015. *Data Deprivation : Another Deprivation to End*. Policy Research Working Paper; no. WPS 7252. Washington, D.C. : World Bank Group. http://documents.worldbank.org/curated/en/2015/04/24426058/data-deprivation-another-deprivation-end

Seth, S., & Alkire, S. (2014). *Measuring and decomposing inequality among the multidimensionally poor using ordinal data: A counting approach*. Queen Elizabeth House, University of Oxford.

Singh, N. and S. Cruz (1997). Issues in local government reform in india. University of California Santa Cruz, Department of Economics Working Paper, November.

Shorrocks, A.F. (1983). Ranking Income Distributions. *Economica* 50: 3-17.

United States Department of Agriculture (2007). Thrifty Food Plan, 2006. http://www.cnpp.usda.gov/sites/default/files/usda_food_plans_cost_of_food/TFP2006Report.pdf

Smith, L. C., O. Dupriez, and N. Troubat, (2014). Assessment Of The Reliability And Relevance Of The Food Data Collected In National Household Consumption And Expenditure Surveys. *International Household Survey Network* working paper.

Solt, F. (2009). Standardizing the World Income Inequality Database. *Social Science Quarterly*, *90*(2), 231-242.

Stiglitz, J. E., A. Sen, and J-F. Fitoussi (2010). *Mismeasuring Our Lives: Why GDP Doesn't Add Up*. New York: New Press.

Smeeding, T., and J.P. Latner, (2015). PovcalNet, WDI and 'All the Ginis': a Critical Review *The Journal of Economic Inequality*, 1-26.

Sukhtankar, S. and M. Vaishnav (2015). Corruption in india: Bridging research evidence and policy options. In India Policy Forum, Volume 11, pp. 193–261.

UNICEF. (1998). The State of the World's Children 1998. New York, UNICEF

Vijaya, R. M., Lahoti, R., & Swaminathan, H. (2014). Moving from the household to the individual: Multidimensional poverty analysis. *World Development*, *59*, 70-81.

Villaseñor, J., & Arnold, B. C. (1989). Elliptical Lorenz Curves. *Journal of Econometrics*, *40*(2), 327-338.

Weisbrod, J., S. Vollmer, Sebastian and H. Holzmann (2007). Perspectives on the World Income Distribution: Beyond Twin Peaks Towards Welfare Conclusions.

World Bank (2014). Summary of results and findings of the 2011 international comparisons program. International Comparison Program, http://siteresources.worldbank.org/ICPINT/Resources/270056- 1183395201801/Summary-of-Results-and-Findings-of-the-2011-InternationalComparison-Program.pdf

World Bank. (2011). *World development report 2011: Conflict, security, and development*. World Bank.

Declaration according to §16 (Assurances) Examination Regulations for the doctoral programme in Economic Sciences

1. The opportunity for the existing doctoral project was not made commercially available to me. Especially, I have not engaged any organisation that seeks thesis advisors against a fee for the preparation of dissertations or performs my obligations with respect to examination components entirely or partly.
2. I declare that I have prepared the submitted dissertation (title follows) independently and without prohibited aids; I have not accepted external help either free-of –charge or against a fee and will maintain this also in the future. I did not make use of any aids and papers other than those indicated by me. I have marked all word-by-word (direct) or implied citations of the writings by other authors.
3. I will adhere to the guidelines to ensure good scientific practice at the University of Göttingen.
4. No equivalent doctoral studies have been applied for at a different university in Germany or abroad; the dissertation submitted or parts thereof have not been used in any other doctoral project.
5. Furthermore, I am aware of the fact that untruthfulness with respect to the above declaration repeals the admission to complete the doctoral studies and/or subsequently entitle termination of the doctoral process or withdrawal of the attained title.

_____

Date, Signature

Declaration about Co-authored Papers

I worked jointly with Sanjay Reddy and Arjun Jayadev on the papers included as the first two chapters of my dissertation. The papers were conceptualized jointly by all three of us. I did the majority of the empirical work and we all contributed to writing and data collection process.

I worked jointly with Stephan Klasen on the paper included as the third chapter of my dissertation. We conceptualized the paper jointly. I did a major portion of the empirical analysis and both of us contributed to the writing of the draft.

I worked jointly with Soham Sahoo on the paper included as the last chapter of my dissertation. We conceptualized the paper jointly. Both of us contributed to the empricial analysis and the writing of the draft.

Date, Signature