

Revisiting Global Poverty Reduction: Public Goods and the World Distribution of Income, 1980-2022

Amory Gethin*

5th January, 2024

Abstract

This article constructs new estimates of global poverty that incorporate the consumption of public services. Combining data from multiple sources, I build a novel historical database on the value and progressivity of public education, healthcare, and other in-kind transfers received worldwide since 1980. Public goods are large and have considerably grown: they represent 30% of global GDP and have been a major driver of inclusive growth. The consumption of public goods accounts for about 20% of global poverty reduction since 1980. Total government redistribution, including cash and in-kind transfers, accounts for 30%. In a companion paper, I incorporate in this analysis the causal impact of education on pretax incomes. Combining indirect investment benefits from education with measures of direct government redistribution brings the contribution of public policies to extreme poverty reduction to at least 50%.

*Paris School of Economics (Grant ANR-17-EURE-001). I am particularly grateful to Thomas Piketty, Vincent Pons, Emmanuel Saez, and Gabriel Zucman for their continuous guidance. I thank Daron Acemoglu, Lars Harhoff Andersen, Alan Auerbach, Pierre Bachas, Thomas Blanchet, Denis Cogneau, Léo Czajka, Edgard Dewitte, Matthew Fisher-Post, Ignacio Flores, Jonathan Gruber, Anders Jensen, Lawrence Katz, Christoph Lakner, Wouter Leenders, Ségal Le Guern Herry, Rosanne Logeart, Clément Mazet-Sonilhac, Marc Morgan, Rowaida Moshrif, Oda Nedregård, Paul Novosad, Tom Raster, Anne-Sophie Robilliard, Justin Sandefur, Todd Schoellman, Camille Semelet, Andrei Shleifer, Morten Nyborg Støstad, David Thesmar, Tom Vogl, Torsten Figueiredo Walter, Duncan Webb, and Ekaterina Zhuravskaya, as well as seminar and conference participants at the Paris School of Economics, UC Berkeley, and the University of Washington for useful comments and discussions.

1. Introduction

Government redistribution is rising around the world. Between 1980 and 2019, real government expenditure per world citizen doubled, from about \$2500 to \$5000 at purchasing power parity. Cash transfers cannot be held responsible: they represent less than 10% of global public expenditure and have scarcely increased since 1980. Instead, the bulk of the growth of government redistribution has been driven by investments in public education, healthcare, housing, police services, transport infrastructure, and other public goods. Together, these transfers represented some 30% of global GDP in 2019.¹

This dramatic transformation remains largely absent from poverty and inequality statistics. The standard concept used to measure global poverty is household final consumption expenditure, defined as the market value of all goods and services purchased by households. By construction, it excludes public goods, since these goods are not bought on a market. As a result, it remains difficult to understand how macroeconomic growth reduces poverty, in a world where almost a third of global GDP is redistributed by governments in unaccounted ways. It also limits our ability to answer some of the most basic questions of human development, such as: who benefits from public goods? How does the provision of public services vary across time and space? And to what extent have public goods contributed to global poverty reduction in the past decades?

This paper represents a first attempt at answering these questions. I propose a simple framework for studying the distribution of public goods that combines two parameters: a cost parameter and a progressivity parameter. The cost parameter corresponds to how much governments spend on each type of transfer. The progressivity parameter governs the share of this transfer that is received by different income groups. I also investigate the robustness of my results to accounting for a productivity parameter, capturing the fact that holding cost constant, the quality of public goods provided may vary across countries, over time, and throughout the income distribution.

I apply this framework to the study of global poverty reduction since 1980. The starting point is a new database on the world distribution of public spending, which I construct by combining data from about twenty different sources. To cover the cost component, I draw on budget data to build new aggregate series on the level and composition of general government expenditure. To cover the progressivity component, I rely on estimates of the distributional incidence of public education and healthcare from various fiscal incidence studies and surveys. To account for potential variations in productivity, I construct measures of cost efficiency by benchmarking

¹See figure 4, which plots the evolution of real worldwide government expenditure per capita since 1980. Appendix figure A.2.25 plots global government expenditure as a share of global GDP.

the value of in-kind transfers to government performance measures. The resulting dataset yields new estimates of the monetary value of public goods received by income group in most countries in the world from 1980 to 2019. It also covers the distribution of taxes and cash transfers, allowing me to compare the incidence of public services to that of these other traditional redistributive tools.

I find that the rising consumption of public goods has played a major role in improving the living conditions of the world's poorest individuals. Figure 1 plots the evolution of the global poverty headcount ratio since 1980 before and after accounting for taxes and transfers. The share of the world's population living with less than \$2.15 per day in 2017 PPP US dollars declined from 23% to 13% in terms of pretax income, representing a 43% decline. After deducting taxes from individual incomes and adding cash and in-kind transfers, this figure rises to 63%. By this measure, government redistribution accounts for about 30% of global poverty reduction since 1980. Public goods alone account for 20%.

Public goods have also played a key role in making global economic growth more inclusive. Public goods tend to strongly reduce inequality within countries, because they are almost always more equally distributed than pretax incomes. As a result, increasing spending on public education, healthcare, and other public services has strongly reduced global inequalities. All income groups within the bottom 60% of the world distribution of income have benefited from greater net government transfers since 1980. The global top 10% to bottom 50% income ratio has declined by 30% before accounting for public services, compared to 36% after doing so. Public goods thus explain almost 20% of total global inequality reduction since 1980. Today, they reduce global income disparities as much as taxes and cash transfers combined.

I also find that dimensions of government redistribution are correlated across countries. In particular, low-income countries score lower on most dimensions of government redistribution. Not only do they spend less on public services, they also invest more heavily in services that are more regressive and provide each of them more unequally. There is also evidence that they provide public services less efficiently than high-income countries, even after accounting for differences in cost of provision. This "triple curse" comes with extreme inequalities in the quality of public services received worldwide. In 2019, only about 0.5% of global GDP was redistributed to the poorest 10% of world citizens, while almost 10% of global GDP accrued to the richest global income decile. As a result, accounting for public goods increases the share of global income disparities explained by inequalities between countries. The share of the global poor living in poor countries is greater than we thought, because citizens of poor countries benefit from public services of much lower quality than those of the rich world.

Together, these results highlight the critical role played by public-private complementarities

in reducing poverty. Economic growth not only improves the labor market and consumption opportunities of low-income households. It also comes with greater government revenue through taxation, a significant fraction of which ends up being redistributed in the form of improved public services. In directly accounting for public goods consumption in the measurement of poverty, my results thus uncover and quantify an important channel—enhanced public spending—through which economic growth contributes to global poverty reduction. It is also important to stress that investments in education, healthcare, and other public services are likely to have also contributed to pretax income growth, in addition to their direct effects on posttax inequality. Accounting for this indirect channel would lead to putting even more weight on public goods in explaining global poverty reduction. In a companion paper, I show that private returns to schooling alone can account for over half of real pretax income gains for the world’s poorest 20% individuals since 1980 (Gethin, 2023a).

Despite their relative robustness, two important limitations of these findings should be acknowledged. A first limitation is empirical: due to the lack of comprehensive data, our understanding of the incidence of many public goods remains quite limited. The approach I adopt thus consists in deriving *lower bounds* on the progressivity and productivity of government expenditure. For instance, I distribute spending on a number of public services proportionally to posttax disposable income, which amounts to assuming that high-income groups benefit from substantially higher transfers. I also make the conservative assumption that the productivity of governments is never higher than that of the private sector. My results can be easily updated and improved as better data becomes available, with the likely conclusion that public goods have contributed to the decline of global poverty and inequality to an even greater extent than estimated here.

A second limitation is more conceptual in nature. While the results presented here provide useful information on the distribution of public goods, they tell us little of their value from the perspective of economic welfare. A classic result of economic theory states that the value of an in-kind transfer should not be higher than that of cash, because cash allows consumers to choose what they consume (Atkinson and Stiglitz, 1976). Yet, a growing literature questions the validity of this claim. For instance, in-kind transfers may be preferable to cash if they insure households against commodity price risk (Gadenne et al., 2022), have larger spillover effects onto children (Hendren and Sprung-Keyser, 2020), or if recipients have a desire for self-control mechanisms (Liscow and Pershing, 2022). There is also survey evidence that individuals may prefer public goods to cash, in particular education and health, both in rich and poor countries (Khemani, Habyarimana, and Nooruddin, 2019; Thesmar and Landier, 2022). I do not attempt to disentangle these different factors here. Put simply, this paper studies the incidence of public services on the distribution of *total consumption*, including both privately and publicly provided

goods, in the same way as GDP is used to compare total production across countries and over time. Moving from consumption-production to economic welfare would require estimating individuals' willingness to pay for the private and public goods that they consume. I discuss challenges in doing so and avenues for future research in this direction in section 5.3.²

This article contributes to our understanding of the evolution of global poverty in the past decades. The classic approach to measuring monetary poverty is to compute the share of individuals whose consumption falls below a given threshold (e.g., [Chen and Ravallion, 2010](#); [Deaton, 2010](#); [Ravallion, 2012](#)). While such measures provide invaluable information on the living standards of the poor, they fail to capture dimensions of economic well-being that are not typically bought on a market. Well aware of this limitation, international organizations and statistical institutes have started developing a number of indicators of multidimensional poverty.³ These different measures have provided useful insights, yet they tend to suffer from limited space and time coverage and are not directly comparable with growth statistics. In this paper, I tackle some of these limitations by constructing measures of monetary poverty and inequality that incorporate public services. I provide evidence that doing so contributes to reconciling monetary and multidimensional approaches to measuring living standards, precisely because public services are major determinants of cross-country differences in deprivation in health, education, and other non-monetary dimensions of quality of life.

This article also provides new evidence on the evolution of global income inequality. A number of studies have attempted to estimate the world distribution of income, generally focusing on household consumption or pretax income (e.g., [Bourguignon and Morrisson, 2002](#); [Chancel and Piketty, 2021](#); [Lakner and Milanovic, 2016](#); [Sala-i-Martin, 2006](#)). I contribute to these efforts by estimating the incidence of all types of taxes and transfers on global poverty and inequality. To the best of my knowledge, this study is the first to analyze how government redistribution in its various forms has contributed to shaping global income disparities since the 1980s.

My methodology is directly inspired by the growing literature attempting to bridge gaps between micro- and macro-approaches to the measurement of living standards. [Piketty, Saez,](#)

²The results presented in this paper can be interpreted as mirroring economic welfare if willingness to pay is exactly equal across all types of private and public goods. It is also important to mention that standard poverty statistics do already incorporate a number of in-kind incomes that are not necessarily optimally “chosen.” These include, for instance, own consumption of food produced by the household and gifts received in kind from other households, both of which may be valued significantly less than cash.

³Such measures have become increasingly available and mobilized in both developed and developing countries: by 2017, 16 countries used multidimensional poverty indices as official measures of poverty ([Glassman, 2019](#)). Since 2010, the Oxford Poverty and Human Development Initiative has published cross-country measures of multidimensional poverty that combine indicators on deprivations in health, education, and living standards ([Alkire, Kanagaratnam, and Suppa, 2021](#)). In the same spirit, the World Bank has recently released a multidimensional poverty measure that incorporates both monetary and non-monetary components ([World Bank, 2018](#)).

and Zucman (2018) construct Distributional National Accounts (DINA) for the United States, allocating the entirety of national income, taxes, and government expenditure to individuals every year since 1913. A number of studies following this framework have been conducted on other countries since then.⁴ The major advantage of this methodology is that it produces estimates of income inequality that are consistent with macroeconomic growth. Its main limitation is that it does not generally account for the progressivity and productivity of public goods. Instead, studies typically assume that all public goods are valued at cost, and received either proportionally to posttax disposable income or as a lump sum.⁵ In this article, I go beyond these simplifying assumptions by explicitly accounting for the progressivity and productivity of public education and healthcare in a national accounts framework.

More generally, this paper extends our knowledge of who benefits from in-kind transfers. A large body of literature has attempted to estimate the distributional incidence of specific public services in specific contexts.⁶ While many of the methods used in this article are directly inspired from this work, I depart from existing studies in taking a long-run, historical perspective on the incidence of all forms of government redistribution on global poverty.

The rest of the paper is organized as follows. Section 2 presents motivating evidence and the general framework used to study the distribution of public goods. Section 3 applies this framework to build a new database on public goods provision worldwide since 1980. Section 4 presents the results. Section 5 investigates the role of potential differences in public sector productivity, discusses how public services can help solving well-known discrepancies between surveys and national accounts in the measurement of poverty, and provides a general discussion. Section 6 concludes.

⁴See in particular Blanchet, Chancel, and Gethin (2022) on Europe, Bozio et al. (2022) on France, and De Rosa, Flores, and Morgan (2022) on Latin America. See also Germain et al. (2021), Bruil et al. (2022), and Jestl and List (2022), who cover posttax income for a limited number of years in France, the Netherlands, and Austria, respectively. See Chancel et al. (2022b) for a presentation of other studies following the DINA methodology.

⁵For instance, Piketty, Saez, and Zucman (2018) allocate all non-health expenditure proportionally to posttax disposable income. Blanchet, Chancel, and Gethin (2022) consider two polar scenarios, one in which public goods are distributed proportionally to posttax disposable income, and one as a lump sum.

⁶See for instance Benhenda (2019), Lustig (2018), Paulus, Sutherland, and Tsakloglou (2010), Verbist, Förster, and Vaalavuo (2012), and Wagstaff et al. (2014) on education and health, Aaberge et al. (2010), 2019 on local government services, and Mladenka and Hill (1978) on police expenditure. To the best of my knowledge, O'Dea and Preston (2010) represents the only attempt at conceptualizing and providing guidelines on how all public services could be allocated to individuals (although they do not attempt to actually do so). My approach is largely inspired by theirs, and in many cases directly follows their recommendations.

2. Motivating Evidence and Conceptual Framework

This section presents motivating evidence for studying the distribution of public goods (section 2.1) and introduces the general framework used in the paper (section 2.2).

2.1. Motivating Evidence

I start by providing motivating evidence for incorporating estimates of public goods delivery in poverty and inequality statistics. I establish two simple stylized facts. First, public and private goods are substitutes: in countries with lower public goods provision, households tend to rely on market alternatives to a greater extent. Second, public goods have large effects on dimensions of well-being that are not captured by private consumption. As a result, standard poverty statistics underestimate poverty in countries with small welfare states relative to those with higher public goods provision. They also tend to structurally underestimate the growth elasticity of poverty, given that economic growth allows governments to invest in public goods that are not recorded in private consumption.

2.1.1. Public and Private Goods Are Substitutes

The standard approach to measuring poverty and inequality focuses on household disposable income or household final consumption expenditure (disposable income minus savings). Disposable income is equal to the sum of labor and capital incomes, minus direct taxes paid, plus cash transfers received. By definition, it excludes public services, which amounts to implicitly assuming that their value to households is exactly zero.

This assumption can lead to implausible conclusions when analyzing the incidence of public policies on poverty. Consider for instance a government that decides to fully subsidize healthcare, effectively bringing down all private out-of-pocket healthcare expenditure to zero. Theoretically, individual incomes should be adjusted by adding the corresponding new in-kind transfer received by the government to their incomes. Yet, in the standard framework, poverty will remain unchanged, because the value of subsidized healthcare is recorded as being exactly zero. More generally, every policy subsidizing the provision of a good that was previously privately bought will be measured as having no incidence on poverty or inequality.

Figure 2a provides evidence that this channel is empirically relevant and quantitatively important. There is a strong negative correlation between the share of households pushed into extreme poverty by out-of-pocket healthcare expenditure and the size of public health spending across

countries. In Bangladesh, where the government spends less than 0.5% of national income on health, 7% of the population see their daily expenditure fall below PPP \$3.65 per day because of private health spending. Meanwhile, less than 0.3% of the South African population ends up poor because of out-of-pocket health spending, in a country where almost 7% of national income is spent on government-provided health services. Private and public expenditure are therefore not independent. In-kind transfers do allow poor households to save money, and not accounting for such money leads to overestimating poverty in countries with large welfare states.

2.1.2. Public Goods Matter for Non-Monetary Dimensions of Quality of Life

Public goods do not only matter for private consumption: they also contribute to improving non-monetary dimensions of well-being. The need to go beyond strictly monetary measures of poverty has been increasingly recognized in the past decades. Accordingly, researchers and international organizations have started developing a number of indicators of multidimensional poverty, which typically involve aggregating individual-level measures of well-being across a number of domains. For instance, [Alkire, Kanagaratnam, and Suppa \(2021\)](#) combine measures of deprivation in health, education, and access to a number of basic goods, each of which is assigned a weight of one-third.⁷

Figure 2b provides suggestive evidence that accounting for in-kind transfers contributes to bridging the gap between monetary and multidimensional poverty statistics. The x-axis plots general government expenditure on education, health, and housing and community amenities as a fraction of net national income. The y-axis represents the difference between the share of households living in multidimensional poverty and the share of households living in monetary poverty. There is a strong negative correlation between the two variables: multidimensional poverty is lower than monetary poverty in countries with large welfare states, while it is significantly higher in countries with low government expenditure. This suggests that in-kind transfers strongly improve the well-being of the global poor in dimensions of quality of life that are not captured by monetary poverty statistics.

The framework adopted in this paper can thus be viewed as one way of incorporating non-monetary dimensions of poverty in a monetary framework, through the value of the public services that largely determine them. The major advantage of this approach is its conceptual consistency with macroeconomic statistics. Unlike multidimensional measures of poverty, it is

⁷More precisely, the index is constructed by attributing a weight of 1/3 to two health indicators (nutrition and child mortality), 1/3 to two education indicators (years of schooling and school attendance), and 1/3 to six “living standards” indicators (access to cooking fuel, sanitation, drinking water, electricity, housing, and basic assets.)

based on an internationally agreed upon framework, the system of national accounts, which remains the most commonly used source for tracking incomes across countries and over time. Unlike classic monetary poverty measures, it accounts for all forms of government spending, which ensures that income estimates incorporate the large fraction of national incomes that is redistributed in the form of public goods.

2.2. Conceptual Framework

I propose to value public goods by combining data on their cost and their incidence throughout the income distribution. Consider individual i receiving pretax labor and capital income m_i , paying taxes $\tau(m_i)$, and receiving cash and in-kind transfers from the government $g(m_i)$. Her posttax income is:

$$y_i = m_i - \tau(m_i) + g(m_i) \quad (1)$$

The value of public goods received is defined as:

$$g(m_i) = \sum_j G^j \times \gamma^j(m_i) \quad (2)$$

G^j is a *cost* component equal to total government expenditure on function j (e.g., education). $\gamma^j(m_i)$ is a *progressivity* component equal to the share of expenditure on function j received by individual i . By definition, $\gamma^j(m_i) \in [0, 1]$.

My benchmark estimates thus amount to valuing cash and in-kind transfers equally, in line with the approach adopted by the national accounts and the existing fiscal incidence literature (e.g., [Lustig, 2018](#); [Piketty, Saez, and Zucman, 2018](#)). A natural concern is that governments may differ in their ability to provide public goods even after accounting for differences in cost of provision. I thus investigate the potential role played by differences in cost efficiency by introducing a third parameter into the estimation:

$$g(m_i) = \sum_j G^j \times \gamma^j(m_i) \times \theta^j(m_i) \quad (3)$$

With $\theta^j(m_i)$ a *productivity* component adjusting expenditure received by i for the quality of the service provided. It equals zero if the transfer is completely useless (for instance, if the value added of teachers at the school attended by i is exactly 0). On the contrary, it may be greater than one if the government is more efficient than a benchmark production unit at providing a

given service (for instance, if public schools are more cost-efficient than private schools). Hence, $\theta^j(m_i) \in [0, +\infty)$, and $\theta^j(m_i) = 1$ corresponds to the case in which public goods are valued at cost of provision.

Given difficulties at conceptualizing and measuring productivity (which explains why national accounts and GDP growth figures do not generally attempt to do so), I start by presenting results with $\theta^j(m_i) = 1$ in sections 3 and 4. I investigate the robustness of my results to departing from this assumption in section 5.

3. Methodology

I now turn to the methodology used to construct a new database on the provision of public services worldwide. I first cover the distribution of pretax income (section 3.1), followed by the estimation of cost (section 3.2) and progressivity (section 3.3). Table 1 provides summary statistics on the data sources and methodology used to distribute government expenditure.

3.1. Pretax Income

The starting point of the construction of the database consists in measuring the distribution of pretax income. Data on global pretax income inequality come from the World Inequality Database ([Chancel and Piketty, 2021](#)), which draws from studies combining surveys, tax, and national accounts data from various sources to build a new database on the distribution of income in all countries in the world since 1980. Average income in each country-year is scaled up to match net national income per capita: poverty and inequality statistics are consistent with macroeconomic growth rates. The concept of income observed is pretax national income, that is, income before accounting for the operation of the tax-and-transfer system, but after accounting for the operation of the pension and unemployment systems.

3.2. Cost G^j

The first step required to distribute public goods is to measure how much governments spend and on which types of policies. To do so, I build a new database on the level and composition of general government expenditure since 1980 by combining various data sources. My primary source for total expenditure as a share of GDP is [Mauro et al. \(2015\)](#), which I complement with other series from the IMF and the IFPRI-SPEED database ([Yu, Magalhaes, and Benin, 2015](#)). For the composition of public spending, I primarily rely on IMF series, which breakdown government

expenditure by Classification of the Functions of Government (COFOG). I combine them with additional data on education, health, and social protection spending from the World Bank, the OECD, and the United Nations Economic Commission for Latin America and the Caribbean.

3.3. Progressivity $\gamma^j(m_i)$

3.3.1. Allocation Principles

Measuring the progressivity of public goods is conceptually and empirically challenging, given that their ultimate beneficiaries cannot always be unambiguously identified. I rely on two key allocation principles to estimate the distributional incidence of public goods, which directly follow the existing literature (e.g., [Lustig, 2018](#); [O'Dea and Preston, 2010](#)). First, public services accrue to individuals based on who receives them at a given point in time. Second, public goods benefit households based on the price they would have to pay to benefit from this service if it was not provided as a public good. These two principles are necessary to ensure conceptual consistency with standard poverty and inequality statistics.

1) Cash Flow Principle First, I distribute public goods to individuals based on their beneficiaries at a given point in time. For instance, education spending is distributed to households who send their children to school, while health spending is distributed to individuals using more intensively the public healthcare system. This ensures that public goods are valued in a way that is conceptually consistent with standard fiscal incidence analysis, which focuses on the incidence of taxes and transfers over a given period. Put differently, public services are allocated in the same way as they would theoretically be if households were to receive a cash transfer at time t and immediately use it to buy the corresponding service on a private market. Departing from this assumption would require moving away from the cross-sectional analysis that forms the basis of international poverty statistics. For instance, high-income earners may benefit from greater public education spending during their lifetime because of longer studies, which implies that education expenditure might be more unequally distributed than generally thought (although only modestly so: see [Riedel and Holger, 2022](#)). Yet, allocating education in this way would also conceptually require moving from the analysis of current income to that of permanent income, incorporating estimates of how much taxes individuals pay over their lifetime and how much cash transfers they receive. Unfortunately, available data does not allow for such a detailed analysis when studying the evolution of global poverty.

2) Equivalent Pricing Principle Second, public goods accrue to households based on the price that they *would have to pay* for the corresponding service, rather than the price they *would be willing to pay*. This ensures again that cash transfers and public goods are valued in a conceptually comparable way: if the household was to receive cash instead of the public good, it would have to pay the market price of the corresponding service to benefit from it, not the maximum value it would be willing to pay. Moving from income to welfare would require accounting for the unobserved value that consumers put on *both* market and public goods. Willingness to pay is higher than the observed price for all consumers located to the left of the demand curve, who would continue buying the good if its price was to marginally increase (e.g., [Aaberge et al., 2019](#)).

In line with standard poverty statistics, which focus on consumption and do not attempt to estimate the individual welfare value of each good bought by each household, I will thus distribute public goods based on who benefits more from them, rather than who might put greater or lower value on each type of service. For example, the welfare perspective would imply that high-income households might be willing to pay significantly more for education, because the real income gains that they would get from returns to schooling might be higher.⁸ This would call for putting a greater value on each dollar of public education received by children from high-income parents. In contrast, assuming that the cost of providing education is the same across income groups, the income perspective implies that education should benefit households proportionally to the number of children attending school. Consistency with standard consumption aggregates thus requires allocating education proportionally to school attendance, not expected real income gains from schooling, because a household willing to send a child to school would have to pay the price of the school, not the price of its returns to schooling, if it was to buy the same service from a private provider.

3.3.2. Education

Public education spending represented about 4.4% of national income in the average country in 2019. Following the existing fiscal incidence literature (e.g., [Lustig, 2018](#)), I distribute education expenditure to individuals proportionally to school attendance of children in the household. The data source is a unique historical micro-database that I have contributed to construct in a companion paper in collaboration with the World Bank ([Gethin, Kofi Tetteh Baah, and Lakner, 2023](#)). The database consists of over 1,300 nationally representative surveys fielded

⁸If the return to schooling is proportional and constant (e.g., 10%) and children from high-income parents can expect to have greater income regardless of education, for instance, then the real expected gains from schooling will be higher from children from high-income parents than those from low-income parents.

in 155 countries from 1980 to 2021. It records detailed information on the structure of the household, school attendance, age, and total household income (or consumption). Based on this information, [Gethin, Kofi Tetteh Baah, and Lakner \(2023\)](#) provide detailed indicators of inequality in access to education and intensity of use of the education system by per-capita household income decile and age in each country. Drawing on this database, I calculate the transfer received by decile d in country c at time t as:

$$G_{dct}^{\text{educ}} = n_{dct}^{\text{pri}} g_{dct}^{\text{pri}} + n_{dct}^{\text{sec}} g_{dct}^{\text{sec}} + n_{dct}^{\text{ter}} g_{dct}^{\text{ter}}$$

Where n_{dct}^j denotes the average number of children in school at level j , g_{dct}^j denotes average spending per child on function j , and $j \in \{\text{pri}, \text{sec}, \text{ter}\}$ refers to primary education, secondary education, and tertiary education, respectively. Data on the relative costs of primary, secondary, and tertiary education per child come from the World Bank's World Development Indicators. The number of children in school by level and per-capita household income is recorded in the [Gethin, Kofi Tetteh Baah, and Lakner \(2023\)](#) database.⁹ Finally, all transfers received in each country-year are proportionally rescaled to match total public education expenditure in the database constructed in section 3.2.

While this approach is straightforward and arguably captures first order differences in access to education, there are two potential sources of concern. The first one is that public education spending may vary not only by level, but also across subnational regions. In particular, poorer regions may benefit from lower spending, leading public education spending inequalities to be underestimated. The second concern is that these estimates do not account for children in private schools, which typically benefit from less (or no) public education spending. This will lead to overestimating public education spending inequalities, since children in private schools tend to disproportionately come from high-income households. I investigate the sensitivity of my results to these two concerns by comparing my estimates to those of the Commitment to Equity Institute (CEQ) Database. The CEQ compiles estimates of tax-and-transfer progressivity from a number of fiscal incidence studies following a comparable methodology (see [Lustig, 2018](#)). Education spending is allocated in the exact same way as above, except that these more detailed studies do exclude children in private schools when allocating transfers and generally also account for variations in spending by subnational region. The CEQ database provides this indicator for one or two years in 45 countries.

A comparison of the two datasets is displayed in appendix figure [A.1.2](#), focusing on the share

⁹When missing, relative costs are assumed to have remained constant before or after the last year available. In the absence of detailed information on school attendance by grade in the microdata, individuals in primary school are taken as those aged 6 to 12, individuals in secondary school as those aged 12 to 18, and individuals in tertiary education as those aged above 18.

of public education spending received by the poorest 50%. The two estimates are strongly correlated, suggesting that the simplified methodology does succeed at capturing broad cross-country variations in education spending inequalities similar to those found by the CEQ. On average, my measures of the bottom 50% share of education spending are slightly lower, mainly because I do not exclude children in private schools from the allocation, while CEQ studies generally assume that they benefit from no public education subsidy at all.¹⁰ This provides reassuring evidence that my estimates provide a good approximation, and if anything likely yield a lower bound on public education transfers received by the global poor.

3.3.3. Health

I distribute health expenditure (3.5% of NNI) proportionally to use intensity of the public healthcare system. Here, I rely directly on the CEQ database, which provides estimates of the distributional incidence of health expenditure from a number of studies. These estimates are typically constructed by using survey microdata covering indicators of frequency of use of public healthcare, such as the number of visits to a public health institution in the past month, or the total amount of user fees paid. These indicators are then aggregated at the household level to derive measures of healthcare use intensity by pretax income decile. The data cover 45 countries for one or two years in the 2010s.

3.3.4. Other Public Goods

Other expenditure includes spending on public order and safety (2% of NNI), transport and other economic affairs (5.8% of NNI), general public services (5.5% of NNI), and defense, housing and community, recreation and culture, and environmental protection (4.6%). In the absence of data on their distributional incidence, I make the conservative assumption that they are received by individuals proportionally to posttax disposable income, that is, in a highly unequal way. I view this as a lower bound. Indeed, there is a case for allocating some of these public services in a much more equal way: for instance, police services can be thought of benefiting households proportionally to the crimes that they experience (e.g., [O'Dea and Preston, 2010](#)), while housing policies include many public housing programs that disproportionately benefit low-income households. In [Gethin \(2023b\)](#), I provide evidence that under reasonable assumptions, nearly all in-kind transfers are more equally distributed than pretax income in the case of South Africa.

¹⁰In practice, the government does contribute to the funding of private schools in many countries, although it usually provides lower funding than to public schools. The true transfer received thus likely falls in-between.

3.3.5. Other Dimensions of Redistributions: Social Assistance and Taxes

Finally, to have a complete perspective on the role of government redistribution in shaping poverty and inequality, I incorporate in the database estimates of the distributional incidence of social assistance and taxes.

Social Assistance I distribute social assistance expenditure (2.9% of NNI on average in 2019) to beneficiaries of cash transfers and in-kind social benefits. The main data sources are [Piketty, Saez, and Zucman \(2018\)](#) for the United States, [Blanchet, Chancel, and Gethin \(2022\)](#) for European countries, the CEQ database (40 countries), and the World Bank’s ASPIRE database (108 countries). In each case, I only distribute social assistance expenditure and exclude pensions and unemployment benefits, given that these transfers are already included in estimates of the pretax income distribution (see section 3.1).

Taxes Finally, I allocate taxes in each country-year by combining data on total tax revenue with estimates of the distributional incidence of taxes. Aggregate data come from [Bachas et al. \(2022\)](#), who build a new database on the level and composition of tax revenue in 150 countries from 1965 to 2018. Data on the share of taxes paid by pretax income decile come from a companion paper ([Fisher-Post and Gethin, 2023](#)).

3.3.6. Imputation of Missing Data

I consider three scenarios for the distribution of public goods, cash transfers, and taxes in countries with missing data. In my benchmark scenario, I fill missing values with the average tax or transfer incidence profile observed in all country-years. I then consider an upper bound in which missing countries are attributed the average incidence profile of the five countries with the most progressive profiles, and a lower bound in which missing countries are attributed the profile of the five countries with the most regressive profiles.

3.3.7. Validation: Comparison With Detailed South African Series

Given the relative scarcity of data, especially when it comes to the time dimension, it is useful to get a sense of how accurately my estimates capture broad trends in government redistribution in countries where more detailed information exists. Appendix figure [A.1.1](#) compares two estimates of the share of national income redistributed to the bottom 50% in the form of public goods in South Africa. The first one corresponds to the “simplified” series estimated

in this paper, which exclusively rely on aggregate budget data from the IMF and the World Bank, estimates of the progressivity of education covering the 2002-2019 period from [Gethin, Kofi Tetteh Baah, and Lakner \(2023\)](#), and estimates of the distribution of health spending for one year from the CEQ database ([Goldman, Woolard, and Jellema, 2020](#)). The second corresponds to “detailed” series constructed in [Gethin \(2023b\)](#). These series combine survey, census, and newly digitized budget data to allocate all public goods to individuals every year since 1993. Unlike simplified series, they cover each function of government in much greater detail, allowing for a precise allocation of local government spending, housing subsidies, public transport, transport infrastructure, police services, and different kinds of subsidies received by households. They cover the evolution of progressivity over time, while simplified series extrapolate the incidence of transfers from one year of data in the case of healthcare. They also account for variations in spending by province, while the simplified series do not.

Despite their limitations, simplified series appear to track remarkably well the evolution of redistribution in South Africa. In both simplified and detailed series, public services received by the bottom 50% are found to have significantly increased over time, from about 7% of national income in 2000 to 10-11% in 2019. If anything, simplified series do slightly underestimate the rise of redistribution, mainly because progressivity is assumed to have remained constant, while [Gethin \(2023b\)](#) finds that it has significantly increased across all functions of government. They also slightly underestimate redistribution in 2019, mainly because housing subsidies and local government expenditure are assumed to be distributed proportionally to posttax disposable income, while [Gethin \(2023b\)](#) finds them to be much more progressive. These results provide reassuring evidence that the simplified allocation developed in this paper provides a very good first-order approximation of levels and trends in government redistribution around the world.

4. Public Goods and the World Distribution of Income

This section presents the main results on the incidence of public goods on poverty and inequality across countries and in the world as a whole. Section [4.1](#) discusses cross-national variations in the size and progressivity of government redistribution around the world since 1980. Section [4.2](#) studies the incidence of public services on global poverty and inequality.

4.1. The Distribution of Public Goods Around the World

I start by exploiting my new database to document three stylized facts on the distribution of public goods. First, public goods are progressive: they systematically reduce inequality. Second,

public goods have grown since 1980, in particular those public goods that are most progressive. Third, redistribution in the form of public goods correlates strongly with economic development: low-income countries spend less on public goods than high-income countries and in ways that are less progressive.

4.1.1. Public Goods are Progressive

Figure 3 plots the distribution of the progressivity of government redistribution across countries, measured as the share of total expenditure received by the bottom 50% (see also table 1). Education, healthcare, and social assistance are all *relatively* progressive (less concentrated than pretax income): they systematically reduce income inequality. However, there are significant variations both across categories and across countries within each category. In particular, cash transfers appear to be *absolutely progressive* in most countries: the bottom 50% receive on average a greater fraction of these transfers than their share in the population. Meanwhile, public education and healthcare tend to be slightly *absolutely regressive*: higher-income earners benefit from greater transfers than low-income groups.

Social assistance is the most progressive function of government, due to the often explicitly pro-poor design of the corresponding programs (such as conditional cash transfers or food stamps). On average, the bottom 50% receives about 64% of social assistance expenditure, However, there are large variations across countries, with the share of social assistance transfers accruing to the bottom 50% ranging from only 16% in Haiti to as much as 92% in Peru.

Education is less progressive than cash transfers but still substantially reduces inequality, as it falls close to a lump sum allocation. In the majority of countries, the bottom 50% benefit from 45-50% of public education expenditure. This figure is the product of two countervailing forces. On the one hand, inequality in access to schooling implies that children from high-income households tend to stay longer in school. The fact that spending per child is higher as higher levels of education reinforces these inequalities. On the other hand, fertility is often slightly higher among low-income households, which increases the progressivity of public education through a demographic effect. These two effects more or less compensate each other on average, yielding a quasi-egalitarian distribution of public education spending.¹¹

Public healthcare is about as progressive as education in the average country, although there are significant variations. In some countries, low-income households use relatively less intensively

¹¹As discussed in section 3.3, in the absence of data, I make the conservative assumption that all children in school benefit from public education spending. If one was to exclude children in private schools, public education would be more progressive, because private schools are used much more intensively by high-income households (Lustig, 2018).

the public healthcare system, partly because user fees may act as a barrier to access. In others, they do so to a greater extent, partly because they suffer from poorer health, and partly because high-income households tend to rely on private healthcare services to a greater extent.

Combining social assistance, education, healthcare, and other public services distributed proportionally to posttax disposable income implies that about 30% of total government expenditure ends up accruing to the bottom 50% in the average country. In nearly all countries in the world, government transfers are relatively progressive (less concentrated than income), but absolutely regressive (accruing in greater proportion to the rich than to the poor). There are large variations in the progressivity of expenditure, with the bottom 50% share of total spending varying from only 15-16% (Angola, Somalia, Republic of the Congo) to almost 50% (Denmark, Sweden, United Kingdom).

4.1.2. Public Goods Have Grown

The second stylized fact is that governments have dedicated growing resources to public services in the past decades. Between 1980 and 2019, average general government expenditure as a share of national income increased from about 26% to 29% (see table 1). This rise cannot be explained by cash transfers: social assistance spending almost stagnated at about 2.5-3% of NNI on average, which represents about 10% of total government expenditure. Much of the rise of government intervention was instead driven by significant increases in public goods, and especially education and healthcare. Meanwhile, expenditure on economic affairs and general public services slightly declined. Overall, net national incomes increased significantly, leading public services to expand considerably in real value in the world as a whole. As shown in figure 4, real government expenditure per world citizen approximately doubled from 1980 to 2019.

Figure 4 breaks down the evolution of government expenditure on social assistance and public services by country income group from 1980 to 2019. There are three main results.

First, low-income countries spend significantly less on both social assistance and public goods as a share of national income than high-income countries. In 2019, total expenditure amounted to about 24% of national income in low-income countries, 27-28% in middle-income countries, and 36% in high-income countries. Poorer countries also dedicate a lower fraction of total expenditure to education and healthcare. Less than 6% of national income is spent on public education and health in low-income countries, compared to almost 15% in high-income countries. Meanwhile, low-income countries actually dedicate a greater share of national income to other public goods than high-income countries (17% versus 15%).

Second, there has been a slight convergence in public goods provision between countries with

different levels of economic development. Total expenditure on public services expanded by about 4 percentage points in low-income countries and 6 percentage points in lower-middle-income countries, compared to about 3 percentage points in high-income countries. It stagnated in upper-middle-income countries, mainly because total expenditure as a share of national income was approximately the same in China in 2019 as in 1980.

Third, there has been a general trend towards devoting greater resources in the most progressive forms of public goods. Regardless of the level of economic development, spending on education and healthcare expanded as a share of national income. In contrast, expenditure on other public goods declined in upper-middle-income and high-income countries and stagnated in lower-middle-income countries. Low-income countries stand out as having invested about as much in education and healthcare as in other public goods.

Combining these results with cross-country differences in macroeconomic growth, middle-income countries appear to have seen expenditure on public services increase most significantly, by about 180% in real terms from 1980 to 2019, mainly due to the rise of China and India.¹² In high-income countries, public goods have expanded almost two times slower, by about 100% from 1980 to 2019. In low-income countries, finally, real expenditure on public goods has almost stagnated, mainly due to exceptionally low or even negative growth in the poorest countries in the world. Differences in public expenditure remain substantial in 2019, with average spending on public goods reaching almost \$15,000 at purchasing power parity in high-income countries, about three times more than in upper-middle-income countries, and over thirty times more than in low-income countries.¹³

4.1.3. Low-Income Countries Score Lower on All Dimensions of Redistribution

The third stylized fact is that there are large variations in redistribution in the form of public goods, which correlate strongly with economic development. Figure 5a maps the share of national income received by the bottom 50% in 2019 around the world. Progressive spending on public goods is generally highest in North America and Western Europe, exceeding more than 8% of national income in most countries. It is also relatively high in South America and Southern Africa, where some countries redistribute similar or even higher shares of national income to the bottom 50% than in Western countries. Public goods provision is significantly

¹²See appendix figure A.2.26, which plots real expenditure on public goods by country income group. Figure A.2.27 plots the same figures expressed as a share of each country's national income.

¹³Appendix figures A.2.28 and A.2.29 plot the corresponding series by world region. Spending on public goods has increased in all regions of the world. This rise has been most pronounced in China and India, and lowest in Africa. See also figure A.5.40, which maps changes in general government expenditure as a share of national income in each country from 1980 to 2019.

lower in Asia: less than 7% of national income is received by the poorest half of the population in most countries. Finally, in-kind redistribution is lowest in Western, Central, and Eastern Africa, where it often falls below 4% of national income.

Figure 5b plots the level and composition of public services received by the bottom 50% in fifteen selected countries or regions, which together represented about two-thirds of the world's population in 2019. There are huge differences in in-kind redistribution to the bottom 50% across countries. In Bangladesh, Nigeria, and Indonesia, only about 4% of national income is received by the poorest half of the population in the form of public goods. The corresponding figure exceeds 11% in Western Europe and the United States. Redistribution in the US is slightly higher than in Western Europe, in line with the findings of [Blanchet, Chancel, and Gethin \(2022\)](#). It is also interesting to note that most differences in in-kind redistribution across countries can be explained by spending on education and health. Less than 1% of national income is received in public education and healthcare by the bottom half of the population in Bangladesh and Nigeria. The corresponding figures are higher than 6% in Brazil, South Africa, Western Europe, and the United States.

While these differences arise from combining data on the size and progressivity of government expenditure in each country, they generally extend to each of these parameters taken separately. Table 2 decomposes the distribution of public goods into these two drivers by country income group and world region. Both dimensions of redistribution increase significantly with economic development. Total expenditure on public goods is about 30% of national income in high-income countries, compared to 23% in low-income countries. 33% of spending accrues to the bottom 50% in the former group, compared to 23% in the latter. Combining these parameters, the bottom 50% ends up benefiting from only 5% of national income in the form of public goods in low-income countries, about two times lower than in high-income countries. Poor countries thus not only invest less in public goods than rich countries; they also provide them much more unequally than in the rich world.¹⁴

Similarly, variations in overall redistribution across geographical regions tend to be reproduced across different dimensions of redistribution. African and Asian countries display significantly lower levels of general government expenditure as a share of national income. They also tend to invest a lower fraction of that expenditure in education and health, the two most progressive

¹⁴Appendix table A.1.1 reports pairwise correlation coefficients between dimensions of redistribution and net national income per capita across countries, including measures of productivity discussed in section 5. Nearly all dimensions of redistribution are significantly positively correlated: countries spending less also spend in more regressive and more inefficient ways. All four parameters are also positively correlated with economic development, in particular progressivity ($\rho = 0.7$) and aggregate productivity ($\rho = 0.65$).

functions of government.¹⁵ China and India stand out as interesting cases. Expenditure on public goods is higher in India, but redistribution is operated in a significantly more progressive way in China. As a result, both countries end up redistributing about 6% of their national incomes to the bottom 50%.

4.2. Public Goods and Global Economic Growth

I now turn to analyzing the incidence of public goods on the distribution of global economic growth since 1980. I first show that public goods have played a major role in making global economic growth more progressive. I then analyze the incidence of public goods on global inequality. Finally, I decompose redistribution into its different components.

4.2.1. Public Goods and Global Poverty Reduction

To what extent has the rise of public goods contributed to the decline of global poverty? Figure 1 plots the evolution of the global poverty headcount ratio at \$2.15 per day, expressed in 2017 PPP USD, before and after accounting for cash transfers and public goods. Following the distributional national accounts methodology (Piketty, Saez, and Zucman, 2018), I compare three concepts of income: pretax national income, posttax disposable income, and posttax national income. Posttax disposable income removes direct taxes from pretax income and adds cash transfers, which corresponds to the standard concept used to measure poverty. Posttax national income removes all taxes, including indirect taxes, and adds all government expenditure, which ensures that average incomes are consistent with net national income growth. Global poverty has declined by about 43% in terms of pretax income, from 23% in 1980 to 13% in 2019. Adding cash transfers lifts about 2% of the world population out of poverty. It also increases the rate of poverty reduction since 1980 to 50%. Finally, adding public goods further reduces poverty by about 4 percentage points, and yields a total rate of global poverty decline of 63%. Hence, government redistribution contributes to reducing the global poverty rate by about a third today, and it has contributed to accelerating the rate of global poverty decline since 1980 by almost 50%. About two-thirds of these effects are driven by public services. Overall, they have contributed to about 20% of the decline in global poverty.¹⁶

¹⁵Online appendix figure A.5.39 maps general government expenditure as a share of NNI around the world in 2019. Figure A.5.41 plots the share of education and health spending in the government budget.

¹⁶It is important to stress that these ratios depend on which transfer is allocated first. If one was to first allocate public goods and then cash transfers, then the contribution of the former would appear substantially higher in comparison to the latter. In allocating public goods after cash transfers throughout the paper, I provide a lower bound on the contribution of in-kind transfers relative to cash transfers.

The key role played by public goods in global poverty reduction can mainly be explained by the rise of public education and healthcare services, which have increasingly accrued to the global poor in the past decades. Figure 6 plots the level and composition of public services received by global bottom 20% since 1980, expressed as a share of total global income. Although redistribution to the bottom 20% in the form of public goods remains extremely low, it has steadily increased in the past decades: about 0.8% of global GDP was redistributed to the global income quintile in 2019, compared to 0.35% in 1980. The bulk of these gains was driven by education and healthcare, whose value was multiplied by three, from about 0.2% to 0.6% of global income. In 2019, they represented over two-thirds of public goods received by the bottom 20%. The combination of increased redistribution with global GDP growth has implied large gains in the real value of public services received by the global poor. From 1980 to 2019, the per capita transfer received by the global bottom 20% was multiplied by about 4, growing from only \$30 to \$120 per year at purchasing power parity.¹⁷

An alternative way of looking at the role of public goods in shaping global poverty reduction is to compare the growth rates of specific groups before and after accounting for public goods. Figure 7 plots the real average income of the world's poorest 20% before and after cash and in-kind transfers. The global bottom 20% average income approximately doubled in terms of pretax income. Adding cash transfers increases this growth rate to over 130%, while incorporating education, health, and other in-kind transfers raises it further to 170%. By this view, cash transfers account for about 25% of global bottom 20% growth, public goods account for 20%, and total transfers account for as much as 40%. Appendix figure A.1.8 extends this analysis to the world's poorest 50% individuals, with similar conclusions.¹⁸

My main result is robust to polar assumptions on the distribution of public goods. On the one hand, one may argue that only education and health eventually accrue to the poor, while other forms of public goods have little value and mostly benefit richer households. On the other hand, there is a case to make for an egalitarian allocation of collective public goods. After all, poorer households do indirectly benefit from services as diverse as street lighting, post offices, environmental protection, local and national administrations, and garbage removal in many countries around the world. Appendix table A.1.2 shows how sensitive is my result on global poverty reduction to these two scenarios. In my benchmark estimates, accounting for public services increases the rate of poverty reduction from 50% to 63%. Restricting public goods to

¹⁷See appendix figure A.1.13, which plots the per capita real value of public services received by the global bottom 20% in 2021 PPP US dollars. Appendix figures A.1.12 and A.1.14 plot the same figures for the global bottom 50%. The results are broadly similar, although education and healthcare represent a slightly smaller fraction of transfers received.

¹⁸More specifically, transfers account for about 20% of real bottom 50% income growth, about 15 points of which is due to public goods and 5% to cash transfers.

education and health leaves this result unchanged. Assuming that all collective public goods are received on a lump sum basis raises the rate of poverty reduction even further, to 79%. My main conclusion is thus relatively robust to different scenarios on the progressivity of other public goods: public services account for 20-30% of global poverty reduction since 1980 and potentially more.

A second concern is that my findings might be driven by a specific country. The obvious candidates are China and India, which together represent over a third of the world's population and have both significantly invested in public services in the past decades. Appendix table [A.1.2](#) reproduces my results on global poverty reduction after excluding China, after excluding India, and after excluding both countries from the sample. The results are qualitatively similar: public services account for about 15% of global poverty reduction when excluding China, 30% when excluding India, and 25% when excluding both countries.

Finally, I investigate the sensitivity of my results to using World Bank data instead of data from the World Inequality Database.¹⁹ The World Bank data cover consumption or posttax disposable income per capita distributions that are not consistent with growth rates reported in the national accounts, so it is not the most adequate data source to study the impact of government redistribution on poverty and inequality. I attempt to reconstruct measures of pretax and posttax income nonetheless, using data available on the World Bank's website.²⁰ The main results are presented in appendix figures [A.1.5](#), [A.1.6](#), and [A.1.7](#) for poverty thresholds at \$2.15, \$3.65, and \$6.85 per day. The results are qualitatively similar to the main findings presented above: redistribution is found to have accelerated global poverty reduction at all thresholds.²¹

¹⁹Both the levels and trends in global poverty in the WID data differ from those of the World Bank for at least four main reasons. First, World Bank estimates focus on consumption (posttax disposable income minus net household saving), while my focus here is on income. Second, the estimates presented here are consistent with national income growth rates, while World Bank estimates are based on surveys and do not attempt to bridge gaps between survey and national accounts aggregates. Third, some of the estimates used in this paper are based on studies relying on data sources that may differ from those of the World Bank in a number of countries, including China ([Piketty, Yang, and Zucman, 2019](#)), India ([Chancel and Piketty, 2019](#)), and Brazil ([Morgan, 2017](#)). See [Chancel and Piketty \(2021\)](#). Fourth, I use GDP purchasing power parity conversion factors, while the World Bank only corrects for price differences in household final consumption expenditure.

²⁰The World Bank does not publish data on the world distribution of income. I thus reconstruct it myself by collecting distributions from the World Bank's website and extrapolating the average income of each country-percentile to missing years using real GDP per capita growth rates. This yields trends in global poverty almost identical to those officially reported by the World Bank. Finally, I reconstruct measures of pretax income as consumption or disposable income, minus cash transfers, plus direct taxes.

²¹Poverty at \$2.15 per day declined by already 77% in terms of pretax income, so it is unsurprisingly difficult to explain much more of poverty reduction with government redistribution. For the two other thresholds, public goods account for a substantial fraction of poverty reduction (about 20% at \$3.65 per day and 30% at \$6.85 per day)

4.2.2. Public Goods and Global Inequality

I now turn to analyzing the incidence of public goods on global income inequality and the distribution of global economic growth.

Figure 8a plots the real income growth rate experienced by each global income percentile from 1980 to 2019. As is well-known (e.g., Chancel and Piketty, 2021; Lakner and Milanovic, 2016), the distribution of global economic growth has taken the shape of an “elephant curve,” being highest at the middle of the global income distribution, lowest for the global upper-middle class, and relatively high among the richest 1%. Yet, little is known of how changes in government redistribution have shaped this general fact. My new database allows for the first time to make progress in answering that question. As shown in figure 8a, the distribution of global income growth has been relatively similar in terms of pretax and posttax disposable income. Higher cash transfers have led to negligible increases in growth rates at the bottom, financed by higher direct taxes paid by global middle- and top-income groups. By this measure, which corresponds to the standard way of studying the incidence of government policies on poverty, redistribution has done little to increase real incomes at the bottom since 1980.

In contrast to cash transfers, public goods have played an important role in making global economic growth more inclusive. The upper line of figure 8a adds public goods to the analysis and removes all taxes so as to reach posttax national income. Moving from posttax disposable income to posttax national income shifts the total growth rate of the 10th percentile from about 100% to 160%. All percentiles within the bottom 60% see their growth rate rise substantially.²² While in terms of posttax disposable income, most percentiles within the bottom 20% grew at a rate lower than that of the top 1%, the opposite is true in terms of posttax national income. Public goods thus appear to have been a major force of inclusive growth since 1980.

Figure 8b represents the evolution of global income inequality since 1980, measured as the ratio of the average income of the top 10% to that of the bottom 50% in the world as a whole, for different income concepts. There are two main results.

First, taxes and transfers significantly reduce global inequality: in 2019, the top 10% to bottom 50% income ratio was 39 in terms of pretax income, compared to 26 in terms of posttax national income. Taxes and transfers all contribute to reducing global inequality, but transfers have the strongest impact. Indeed, cash transfers reduce the indicator by about 5 percentage points;

²²Appendix figure A.1.9 compares the growth rates of average disposable income and of public services received by global income percentile from 1980 to 2019. In line with the results presented above, public goods have grown significantly faster than posttax disposable incomes, especially at the bottom of the global income distribution. Total disposable income growth ranges from 80% to 180% within the global bottom 50%, while total growth in public services received ranges from 220% to 360%.

adding in-kind transfers further decreases it by 5 percentage points; finally, removing taxes pushes it down by 4 percentage points. By this measure, transfers account for about 70% of the impact of government redistribution on global inequality, while taxes account for about 30%.

Second, public goods have been the strongest driver of the rise of global government redistribution since 1980. Global pretax income inequality has fallen in the past decades: the richest decile earned 53 times more than the poorest half of the world's population in 1980 compared to 39 times today, amounting to a 26% decline. The corresponding figures are 30% after cash transfers, 37% after cash and in-kind transfers, and 37% after all taxes and transfers. In other words, accounting for government redistribution increases the total decline in global income disparities since 1980 by 40%. About two-thirds of this effect is driven by public goods.²³

Because public goods provision varies so widely across countries, it does not only affect poverty and inequality in the world as a whole: it also shapes their distribution across space.

The upper panel of table 3 provides a Theil decomposition of global inequality into its between-country and within-country components for the main income concepts of interest. In line with the results presented above, taxes and transfers reduce global inequality: the Theil index is 1.13 in terms of pretax income, 0.98 in terms of posttax disposable income (or 13% lower), and 0.8 in terms of posttax national income (or 29% lower). However, because poor countries tend to have less progressive tax-and-transfer systems, and because redistribution only reduces inequality within countries, it increases the share of global income disparities explained by inequality between countries. The between-country component accounts for 30% of global inequality in terms of pretax income, but 33% in terms of posttax disposable income, and as much as 39% in terms of posttax national income. Accounting for government redistribution, in particular public goods, thus increases the weight of national differences in net national incomes per capita in explaining global inequality.²⁴

The lower panel of table 3 focuses more specifically on the bottom of the distribution by breaking down the geographical location of the world's poorest 20% by world region. Accounting for government redistribution significantly increases the share of the global poor living in India, Pakistan, Bangladesh, Ethiopia, Nigeria, and other Sub-Saharan African countries, all of which were identified previously as having weak and regressive tax-and-transfer systems. On the contrary, it improves the relative positions of low-income individuals living in China, Latin

²³Appendix figures A.1.15 and A.1.16 plot the evolution of the Gini and Theil indices of global inequality for different income concepts. The conclusions are qualitatively similar: for both indicators, accounting for public goods leads to a faster decline in global income disparities since 1980.

²⁴Appendix figure A.1.17 plots the share of global inequality explained by average income differences between countries from 1980 to 2019 for different income concepts. This share has significantly declined across all income concepts.

America, and the Western world. These differences are quantitatively large. For instance, moving from pretax income to posttax national income increases the share of the global bottom quintile living in India from 18% to 24%, while this share drops from 7% to almost zero in Western Europe and North America.

In the end, lower redistribution in low-income countries translates into huge inequalities in the quality of public services received around the world. In 2019, public goods benefiting the poorest 10% of the world's citizens represented less than 0.5% of global GDP. The share of global GDP received by global bottom 50% as a whole increased significantly throughout the period, from about 1.5% to 3.5% of global GDP, mainly due to greater education and health transfers. However, this remains extremely small in comparison to the quality of services enjoyed by the richest world citizens: in 2019, public goods received by the upper decile of the global income distribution amounted to over 10% of global GDP.²⁵

4.2.3. Decomposing Redistribution

Figure 9 further breaks down the incidence of government redistribution by showing how the global poverty rate behaves under a number of counterfactual scenarios on the size and progressivity of taxes and transfers.²⁶ The three leftmost bars show that taxes and transfers reduce global poverty from about 13% to 7%, as in figure 1.

The next bar considers a radical scenario in which government expenditure would be distributed on a lump sum basis, that is, in a perfectly egalitarian way ($\gamma^j(m_i) = \gamma$). This would reduce global poverty by 4 percentage points. This large effect is consistent with the significant inequalities in the distribution of public goods documented above and the fact that these inequalities are particularly high in poor countries, which spend less on the types of public goods that are most progressive. The last two bars further impose that all countries in the world move to a “Nordic welfare state,” redistributing 50% of their national income, and that the global poor do not have to pay taxes to finance this expenditure. Moving to a Nordic welfare state would have a large effect on global poverty, while removing taxes would reduce it only marginally. This finding is consistent with the fact that both taxes and transfers are substantially lower in poor countries than in the rich world. Overall, applying all these scenarios jointly would reduce the global poverty rate from about 7% to below 1%.

In summary, about 3-6% of the world's population falls below the poverty line because of inequalities in access to public services. Equalizing the distribution of transfers and increasing

²⁵See appendix figures [A.1.10](#) and [A.1.11](#).

²⁶Appendix figures [A.1.22](#) and [A.1.23](#) extend this analysis to poverty at \$3.65 and \$6.85 per day and also incorporate scenarios on public sector productivity discussed in section [5.1](#). The results are similar.

government capacity would have the biggest incidence on global poverty, followed by improving tax progressivity. Even under extreme scenarios on the size and progressivity of government transfers, however, the global poverty rate would still reach about 1%. This points to the roles of both cross-country macroeconomic convergence and reductions in pretax income inequality within countries as necessary complementary factors for improving the living conditions of the global poor.

5. Discussion and Extensions

This section briefly discusses some implications of the results presented in this article and avenues for future research. Section 5.1 explores the robustness of my results to accounting for public sector productivity. Section 5.2 investigates how accounting for public services can shed new light on a key debate in development economics: whether surveys or national accounts should be used to track poverty and economic development. Section 5.3 discusses challenges in moving from measures of consumption to measures of the welfare value of public goods. Section 5.4 explores the potential of my new measures of public goods redistribution for the study of the political economy of inequality.

5.1. Accounting for Public Sector Productivity

A natural concern is that cost of provision may not be an accurate indicator of the quality of public services received, because the productivity of governments may vary across time and space. In this section, I investigate the robustness of my results to adjusting transfers received for variations in public sector productivity. I focus on the main results and leave an extended presentation of the methodology to appendix B.

5.1.1. Methodology

Conceptual Framework I consider an extension in which the value of public goods is allowed to differ from cost of provision. The value of public goods received by individuals can theoretically be broken down into three components:

$$g(m_i) = \sum_j G^j \times \gamma^j(m_i) \times \theta^j(m_i) \quad (4)$$

With G^j government expenditure and $\gamma^j(m_i)$ the share of expenditure received by i . $\theta^j(m_i)$ captures the fact that for a given cost of provision, individuals may receive services of different quality. Empirically, it is useful to make a distinction between two notions of productivity:

$$\theta^j(m_i) = \Theta^j \times q^j(m_i) \quad (5)$$

Θ^j is the *aggregate productivity* of expenditure on function j , which does not depend on m_i . It captures the fact that the government may be more or less efficient at providing a given service than a benchmark production unit. For instance, public schools in country A may be on average less cost-efficient than public schools in country B, which implies that all public education transfers should be reduced by a constant factor in country A.

$q^j(m_i)$ is a *heterogeneous productivity* parameter. It captures the fact that the quality of services provided, holding cost constant, may differ between income groups. For instance, teachers teaching in poorer areas may be more or less qualified than those teaching in richer areas, independently from the wages they receive.

Aggregate Productivity Θ^j I propose to estimate the productivity of public education and healthcare by anchoring cost of provision to educational and health outcomes. For education, the outcome of interest is expected human capital at age 5, which I derive by combining data on school attendance and test scores from international databases. For health, the outcome is the healthcare access and quality index provided by the global burden of disease study ([GBD, 2022](#)), which ranks healthcare systems from 0 to 100 based on death rates from 32 causes of death that could be avoided by timely and effective medical care. I choose these indicators for two main reasons. First, they are among the only education and health indicators for which data is available for almost all countries in the world and with some time dimension. Second, they are relatively good measures of the output of the public sector, in contrast to other measures such as life expectancy, which are arguably more contaminated by unobserved factors.

I then compare these outcomes to spending on education and healthcare to derive measures of cost efficiency in each country-year. Appendix figures [A.3.32](#) and [A.3.33](#) provide a concrete illustration. Education spending per capita is strongly correlated with expected human capital, but there is also significant variation in educational outcomes for a given level of spending. Country-years that perform best for a given cost are attributed $\Theta^j = 1$: they are at the “efficient frontier”. Meanwhile, country-years below the frontier are attributed lower values of Θ^j the further they are from the frontier. In this approach, no country-year has a score higher than 1, implying that the best government in the world is assumed to do just as well as the private

sector and never better.

I discuss the limitations and implications of this approach in appendix 5.1. I view these estimates of productivity as a lower bound for three reasons. First, PPP conversion factors already make an adjustment for public sector productivity, so this approach holds the risk of “double-counting” inefficiencies (World Bank, 2013). Second, they imply necessarily reducing transfers in all countries that are not at the frontier ($\Theta \leq 1$). This is equivalent to assuming that governments are never more efficient than the private sector: absent any government, education and healthcare would be delivered at the same price or lower in any country-year. Third, omitted variable bias implies that productivity is likely to be underestimated in low-income countries, whose lower educational and health outcomes are arguably the product of other factors than government performance (such as lower income *per se*). That being said, I find that my measures of productivity correlate positively with existing indicators of government efficiency, which I view as reassuring evidence that this approach captures cross-country differences in public sector productivity relatively well.

Heterogeneous Productivity $q^j(m_i)$ Heterogeneous productivity is arguably even more challenging to estimate. In the absence of better data, I investigate using subjective perceptions of public services from international survey data to derive estimates of heterogeneous productivity by income group around the world. The data source is the Gallup World Poll, a yearly survey conducted since 2005 in 165 countries, which asks respondents whether they are satisfied with different types of public services in their area. I aggregate average responses by income quintile to measure differences in satisfaction with local public education, healthcare, police, and transport services. I then use relative responses as a scaling parameter, to increase or decrease the transfer received by each income group, for each of these four functions of government. This approach is arguably far from being satisfying. Nonetheless, existing empirical evidence on inequalities in service delivery (conditional on access) suggest that heterogeneous productivity is likely to be quantitatively modest (see appendix B).

5.1.2. Main Results

Cross-Country Differences in Redistribution The main takeaway is that accounting for productivity magnifies cross-country differences in redistribution. Appendix table A.1.3 extends table 2 to productivity-adjusted estimates. Mechanically, because $\Theta^j \leq 1$ by assumption, all countries end up redistributing a lower fraction of national income to the bottom 50%. The gap is particularly large in the case of low-income countries: in-kind transfers are reduced by 40%, compared to about 15% in high-income countries. By this view, poor countries thus suffer

from a “triple curse” of redistribution in the form of public goods: not only do they spend less on public services and distribute them more unequally, they also provide them less efficiently. The gap in total spending between low-income and high-income countries is about 30% (23% versus 30%), while the gap in the value of the transfer eventually accruing to the bottom 50% exceeds 250% (3% versus 8%).

In terms of regional patterns, Asian and Sub-Saharan African countries are characterized by significantly lower aggregate and heterogeneous public sector productivity than Western countries and Latin America.²⁷ The China-India comparison is also striking: although India spends more on public goods than China, the productivity-adjusted transfer received by the bottom 50% ends up being a third lower in India than in China as a share of national income. This finding is consistent with the literature documenting the exceptionally low performance of the Indian public sector (Das et al., 2016; Muralidharan, 2019; Muralidharan and Sundararaman, 2015).

Global Poverty and Inequality I now turn to implications of productivity adjustments for the analysis of global poverty and inequality. The main conclusion is that accounting for productivity substantially reduces in-kind transfers received by the global poor. However, it does not significantly alter the trend; as a result, it only marginally affects my results on the role of public goods in reducing global poverty and inequality.

Appendix figure A.1.18 plots the share of global income received by the world’s poorest 20% before and after adjusting for productivity. Adjusting for productivity reduces the total transfer received by the global bottom 20% by about a third but does not affect the trend. Productivity-adjusted estimates suggest that the share of global income accruing to the poorest quintile rose from about 0.2% to over 0.5%.

Appendix figures A.1.19 and A.1.20 turn to global poverty reduction and the distribution of global economic growth. Adjusting for productivity reduces the rate of global poverty reduction from 63% to 60%. It also reduces the growth rate of percentiles within the bottom 50% of the world distribution of income by 5 to 10 percentage points out of growth rates ranging from 150 to 220. My main findings thus appear to be relatively robust to accounting for a potentially lower productivity of the public sector in low-income countries.

Finally, to get a sense of the importance of productivity in shaping the relationship between

²⁷See online appendix figures A.5.43 and A.5.44, which map aggregate education and health productivity scores in each country. Figure A.5.47 maps average differences in satisfaction with public services across countries. In both dimensions, however, available data suggests that there has been a convergence over time: see figures A.5.45 and A.5.46 for aggregate productivity, and figure A.5.48 for income differences in satisfaction with public services.

public goods and global poverty reduction today, appendix figure [A.1.21](#) reproduces figure 9 with additional steps in which aggregate and heterogeneous productivity differences would be eliminated. With productivity adjustments, the poverty rate in 2019 is about 8% after accounting for all taxes and transfers. Removing heterogeneous productivity differences (setting $q^j(m_i) = 1$) would reduce poverty by less than half a percentage point. Remove aggregate productivity differences (setting $\Theta^j = 1$) would have a larger effect, reducing poverty by about one percentage point. These effects are significant, but still much lower than the effect of equalizing all transfers. These results suggest that improving productivity can be useful to reduce global poverty, but reducing inequalities in access to public services is likely to have quantitatively larger effects.

5.2. Surveys, National Accounts, and Public-Private Complementarities: Public Goods and Measurement Discrepancies in Poverty Statistics

A major debate in development economics centers around whether national accounts or surveys should be used in priority to measure economic development and poverty in the developing world. For reasons that continue to not be well understood, persistent discrepancies between GDP and survey incomes can lead to conflicting conclusions on the evolution of living standards in the past decades ([Deaton, 2005](#)).

Recent studies point to GDP as providing a better benchmark for tracking economic development than household surveys. Combining data from various sources, [Pinkovskiy and Sala-i-Martin \(2016\)](#) provide evidence that GDP correlates much more significantly to satellite-recorded nighttime lights than survey means. It also accounts for a much greater fraction of variations in a number of indicators of quality of life, such as life expectancy, access to safe water, and primary school enrollment. Most importantly, the difference between GDP and survey means is positively associated with achievements on these indicators. In other words, “countries with higher and growing well-being tend to suffer from progressively greater mismeasurement of income by surveys.” While the authors suggest that this finding could be due to the complexity of survey questionnaires, the exact reasons underlying this result remain unclear.

There is one natural candidate for explaining this discrepancy: public goods. As was made clear from the results presented in this article, surveys entirely miss services provided by governments in the form of education, health, transport, and other public services, which are not bought on a market and are thus absent from standard consumption measures. Arguably, these services play a key role in improving quality of life in the exact dimensions studied by [Pinkovskiy and Sala-i-Martin \(2016\)](#), as was already suggested in Figure [2b](#). The share of national income

spent on public goods also appears to have significantly risen in the past decades, which could partly explain why surveys and GDP have become increasingly disconnected from each other. I investigate this possibility in appendix table A.1.4. In the spirit of [Pinkovskiy and Sala-i-Martin \(2016\)](#), I regress five indicators of quality of life on the gap between GDP per capita and survey means: expected years of schooling, youth literacy, the secondary school enrollment rate, infant mortality, and life expectancy. I then compare the coefficient obtained before and after controlling for public spending on education and health, taken as a proxy for public goods provision in these two dimensions of well-being.

In line with [Pinkovskiy and Sala-i-Martin \(2016\)](#), I find that the gap between GDP and surveys tends to be positively correlated with greater quality of life, both before and after adding country fixed effects (panels A and B). For instance, a 1% increase in the gap between GDP per capita and average survey income is associated with a 0.16% increase in expected years of schooling. However, controlling for spending on education or health considerably reduces the size of the coefficient and renders it statistically non-significant in most specifications. Put differently, one of the main reasons why GDP estimates track indicators of quality of life better than surveys is that they incorporate consumption of public goods while surveys do not. In directly incorporating this “missing consumption” into poverty and inequality statistics, this article contributes to correcting some of the conceptual discrepancies between these two approaches to the measurement of living standards. My results also highlight the critical role played by public-private complementarities in global poverty reduction: by enhancing public spending possibilities through greater tax revenue, GDP growth allows governments to increase public goods provision. Accounting for this channel, as was done in this article, leads to a more positive view of the role of macroeconomic growth in reducing poverty than the one pictured by household surveys alone.

5.3. From Consumption to Welfare: Challenges in Measuring the Value of Public Goods

A key limitation of the results presented in this article is that they do not account for how “valuable” public services actually are. While receiving free education might be useful to low-income households, it might not be as useful as receiving food or cash. In this section, I briefly discuss conceptual and empirical challenges in estimating the value of public goods.

There are at least three alternative ways of measuring the value of public services: through stated preferences, through revealed preferences, and through outcome-based estimation.

Stated Preferences Stated preferences refer to what households actually consider the value of public goods to be. To the best of my knowledge, only two studies have attempted to explicitly ask households whether they would prefer receiving cash than public services, and in what proportions (Khemani, Habyarimana, and Nooruddin, 2019; Thesmar and Landier, 2022).²⁸ In both cases, public services are found to be preferred to cash by a majority of households, in particular education and health.²⁹ By this measure, at least some public services should be attributed a greater value than cash transfers, which would reinforce my finding on the role of public goods in reducing global poverty.

Revealed Preferences Revealed preferences approaches use various methods to derive implicit measures of households' willingness to pay for public goods from behavioral patterns. The underlying principle is quite simple: if households receiving a cash transfer do not use it entirely to buy more education, then any increase in education spending should be attributed a lower value than a cash transfer of the same amount. This is a classic finding of economic theory (Atkinson and Stiglitz, 1976): cash transfers are superior to in-kind transfers, because they allow households to choose what they consume.³⁰ Based on this general result, public services should be attributed a significantly lower value than cash when being incorporated into poverty statistics, because they are not “freely chosen” by households.

Outcome-Based Measures Finally, outcome-based approaches value public services based on their actual effects. For instance, Finkelstein, Hendren, and Luttmer (2020) propose to measure the value of public policies by comparing the cost of each policy to total returns for its beneficiaries. In this context, the relative value of public services with respect to cash transfers depends on their ability to improve welfare. Focusing on 133 policy changes in the United States, Hendren and Sprung-Keyser (2020) provide evidence that investments in health and education targeted to low-income children display the highest marginal value of public funds, because they end up paying for themselves through substantial increases in earnings in later life.

²⁸See also Liscow and Pershing (2022), who test the preferences of US citizens for in-kind transfers compared to cash, but focusing on a basket of basic necessities, not on public goods.

²⁹Thesmar and Landier (2022) ask respondents in France, Germany, and the United States to compare the actual composition of the government budget to the one they would prefer. They find clear majority support in favor of greater spending in education and health, and lower spending in cash transfers and defense. Khemani, Habyarimana, and Nooruddin (2019) perform a similar exercise in the context of Bihar, India.

³⁰Another approach consists in using housing prices to derive implicit valuations of public services. For instance, Eshaghnia, Heckman, and Razavi (2021) find, drawing on granular data on housing prices and school characteristics in Denmark, that low- and high-income households are willing to pay a relatively similar fraction of their income for an increase in school quality (see Eshaghnia, Heckman, and Razavi (2021), Figure 4). By this measure, high-income households put a much greater monetary value on education than low-income households, which would imply distributing education spending in a more unequal way than done in this article.

This would call for potentially putting a greater value on education and health expenditure than on cash transfers. Extending this approach to the study of global poverty would ideally require estimating the marginal value of an extra dollar spent in different types of public services in each country. These estimates could then be used to value public services by comparing their marginal value to that of cash transfers.

Understanding Discrepancies In a world with full information, perfect rationality, and perfectly competitive markets, these three measures of the value of public services should coincide, because households would be willing to pay a price equal to expected returns. However, this is rarely the case for at least three reasons.

First, many of the assumptions underlying the Atkinson-Stiglitz theorem do not hold in practice. Poor households may spend little on education and health not because returns are low, but because of many other factors such as limited information on their actual benefits, liquidity constraints, and market imperfections or spatial frictions that limit the supply of private education and healthcare services. All these factors are likely to lead to downward-biased estimates of willingness to pay when measured from revealed preferences.

Second, individuals may value public services beyond the direct value that they get from consuming them. Support for government provision of services is not only dictated by personal benefits, but also strongly responds to beliefs about what constitutes a just society ([Thesmar and Landier, 2022](#)). Public goods may have positive externalities, such as lower inequality, of which individuals are well aware; knowledge of these externalities causally increases support for redistribution ([Lobeck and Støstad, 2022](#)). Outcome-based measures do not generally account for these externalities, which could lead to underestimating the true value of public services.

A third discrepancy comes from the fact that stated preferences may be subject to considerable measurement error, depending on the way questions are framed and other characteristics of survey design. As in the case of revealed preferences, individuals may also not be fully informed about how valuable public services are compared to one another and compared to cash. This makes it difficult to use stated preferences as a benchmark for valuing public goods.

All these inconsistencies make it difficult to evaluate the exact value that should be attributed to public services, both theoretically and empirically. This value ultimately depends on what one believes should matter, whether it is what individuals want (stated preferences), what they actually do (revealed preferences), or the benefits that they eventually get and what kinds of benefits are most important (outcome-based measures). Arguably, all of these three dimensions of welfare matter and should be studied jointly in future research.

5.4. The Correlates of Public Goods Provision: An Exploratory Analysis

I conclude this article with an exploratory analysis of the cross-country correlates of public goods redistribution. The objective is not to provide any new causal evidence, but merely to illustrate how the measures constructed in this article could contribute to shedding new light on the political economy of inequality. I hope that the methodology developed in this article, focusing not only on how much governments spend but also on how progressively and efficiently they do so, can inspire new studies on the different modalities through which public policies can reduce poverty. Combining subnational data on political outcomes with indicators on the size, progressivity, and productivity of public goods provision would be a particularly fruitful avenue for future research.

I investigate the correlates of redistribution in the form of public goods by combining my new measures with selected political and economic indicators available from international datasets. The outcome of interest is the share of national income received by the bottom 50% in the form of public goods in each country, computed from the database constructed in this article. I consider five explanatory variables. The first two capture political regime characteristics: the electoral democracy index available from the V-Dem database and the political competition index produced by the Polity5 project. The next two are measures of public sector corruption (V-Dem) and government effectiveness (World Bank Worldwide Governance Indicators), which relate more closely to the quality of governance. The last variable is the log of GDP per capita, expressed in 2021 PPP US dollars. All models control for the level of inequality, the total population, the demographic structure, and the trade to GDP ratio.

The results of this exercise are presented in appendix table [A.1.5](#). The first three columns correspond to pooled OLS regressions on the full sample (column 1), the 2000-2019 sample only (column 2), and the 2000-2019 sample after excluding advanced Western democracies (column 3). Columns 3 to 6 repeat the same three specifications with country fixed effects.

Pooled OLS regressions point to the electoral democracy index, government effectiveness, and economic development as being significantly associated with redistribution. Electoral democracy and GDP per capita predict greater pro-poor spending on public goods, while government effectiveness has the opposite effect. The latter result might be driven by the fact that more effective governments spend less on public goods because they are able to provide them in more cost-efficient ways. Public sector corruption is associated with lower redistribution, but the effect is smaller and only statistically significant at the 10% level in the second specification.

Electoral democracy stands out as the only robust correlate of redistribution when adding

country fixed effects. This effect is large, statistically significant, and relatively stable across specifications. Moving from the least democratic to the most democratic regime is associated with an increase in public goods received by the bottom 50% of 0.7 to 1.4 percentage points of national income. In contrast, political competition, public sector corruption, government effectiveness, and GDP per capita all display smaller and statistically non-significant coefficients in most specifications.

While these results are only suggestive and should be interpreted with care, they resonate well with the large literature pointing to the key role of political representation in fueling the rise of the welfare state (e.g., [Cascio and Washington, 2014](#); [Fowler, 2013](#); [Fujiwara, 2015](#); [Lindert, 1994](#); [Meltzer and Richard, 1981](#)). They are also in line with recent evidence ruling out the “luxury good hypothesis,” according to which social protection would be a luxury good mechanically growing over the course of economic development ([Lokshin, Ravallion, and Torre, 2022](#)). After controlling for political variables and including country fixed effects, GDP per capita is not significantly associated with more or less redistribution.

6. Conclusion: Three Proposals to Improve Poverty Statistics

Public goods matter. They have been major drivers of human development in the past decades, contributing to improved access to education, healthcare, security, and other dimensions of quality of life. Yet, still little is known of who exactly benefits from these services, not only in a given country but even less so in the world as a whole. This article represented a first attempt at incorporating measures of public service delivery in global poverty statistics. I showed that doing so leads to a more positive view of global poverty reduction since 1980, because public goods are strongly progressive and governments have been increasingly investing in them. Nonetheless, the share of the world’s GDP accruing to the global poor remains extremely limited, because low-income countries suffer from a curse of providing public goods in lower quantities, less progressively, and also potentially less efficiently than in the rich world. There is space for improvement in all three of these dimensions of government redistribution. Enhancing tax revenue, improving equity in access to public services, and raising government productivity should be seen as necessary complementary tools in the fight against global poverty.

This article has taken a large, global perspective on poverty reduction in the past decades, yet much remains to be done to better track public goods delivery around the world. First, there is an urgent need for more transparency on what governments actually do. The data exploited in this article cover spending on large categories, such as education or health, without much detail on the underlying policies. Unfortunately, information on these policies remains very

limited; even when it exists, it often ends up buried under a multitude of documents published by different institutions. The publication of regular reports consolidating and harmonizing data on government budgets, with precise information on the corresponding policies, should be viewed as a priority not only for government accountability, but also for the measurement of global poverty. Too often, researchers and statistical institutes aiming to track living standards face no other option than to ignore public services, simply because of a critical lack of data on what these public services actually are.

Second, more attention should be given to public goods in the design of living standards surveys. Surveys routinely fielded by statistical institutes spend considerable time and effort compiling detailed data on household expenditure, yet the information that they collect on access to basic public goods remains rudimentary at best. Adding regular questions on both objective indicators and subjective perceptions of public service delivery would allow for a much more complete view of the well-being of low-income households. These questions should be designed in ways that would make them directly comparable with spending data on the different kinds of public services provided by governments.

Third, much more research should be conducted on how individuals actually value public services, not only in comparison to cash but also in comparison to one another. Under which conditions do households prefer to receive a transfer in the form of public healthcare, rather than education or cash? How do these priorities vary across countries, over time, and throughout the income distribution? Evidence on these questions remains extraordinarily scarce. Designing surveys eliciting such preferences would represent an important contribution to our understanding of the role of public services in reducing global poverty. Ideally, specific modules could be directly added to the questionnaires of living standards surveys, so as to regularly collect information on citizens' needs and priorities when it comes to public goods delivery.

References

- Aaberge, Rolf, Manudeep Bhuller, Audun Langørgen, and Magne Mogstad (2010). "The distributional impact of public services when needs differ". In: *Journal of Public Economics* 94.1, pp. 9–10.
- Aaberge, Rolf, Lasse Eika, Audun Langørgen, and Magne Mogstad (2019). "Local governments, in-kind transfers, and economic inequality". In: *Journal of Public Economics* 180.1, pp. 1–22.
- Adam, Antonis, Manthos Delis, and Pantelis Kammas (2011). "Public sector efficiency: leveling the playing field between OECD countries". In: *Public Choice* 146.1, pp. 163–183.

- Afonso, António, Ludger Schuknecht, and Vito Tanzi (2005). “Public sector efficiency: an international comparison”. In: *Public Choice* 123.1, pp. 321–347.
- Alkire, Sabina, Usha Kanagaratnam, and Nicolai Suppa (2021). “The Global Multidimensional Poverty Index (MPI) 2021”. Oxford Poverty and Human Development Initiative.
- Altinok, Nadir, Noam Angrist, and Harry A. Patrinos (2018). “Global Data Set on Education Quality (1965-2015)”. World Bank Policy Research Working Paper No. 8314.
- Atkinson, Anthony B. and Joseph E. Stiglitz (1976). “The Design of Tax Structure: Direct Versus Indirect Taxation”. In: *Journal of Public Economics* 6.1, pp. 55–75.
- Bachas, Pierre, Matthew H. Fisher-Post, Anders Jensen, and Gabriel Zucman (2022). “Globalization and Factor Income Taxation”. NBER Working Paper 29819.
- Benhenda, Asma (2019). “Teaching Staff Characteristics and Spending per Student in French Disadvantaged Schools”. Working Paper.
- Blades, Derek (2013). “Health and Education”. In: *Measuring the Real Size of the World Economy: The Framework, Methodology, and Results of the International Comparison Program – ICP*. Ed. by World Bank. The World Bank. Chap. 11, pp. 301–318.
- Blanchet, Thomas and Lucas Chancel (2016). “National Accounts Series Methodology”. World Inequality Lab Working Paper 2016/1.
- Blanchet, Thomas, Lucas Chancel, and Amory Gethin (2022). “Why Is Europe More Equal Than the United States?” In: *American Economic Journal: Applied Economics* 14.4, pp. 1–40.
- Bourguignon, François and Christian Morrisson (2002). “Inequality among World Citizens: 1820-1992”. In: *American Economic Review* 92.4, pp. 727–744.
- Bozio, Antoine, Bertrand Garbinti, Jonathan Goupille-Lebret, Malka Guillot, and Thomas Piketty (2022). “Predistribution vs. Redistribution: Evidence from France and the U.S.” WID.world Working Paper.
- Bruil, Arjan, Céline van Essen, Wouter Leenders, Arjan Lejour, Jan Möhlmann, and Simon Rabaté (2022). “Inequality and Redistribution in the Netherlands”. Working Paper.
- Cascio, Elizabeth U. and Ebonya Washington (2014). “Valuing the Vote: The Redistribution of Voting Rights and State Funds following the Voting Rights Act of 1965”. In: *Quarterly Journal of Economics* 129.1, pp. 379–433.
- Chancel, Lucas, Denis Cogneau, Amory Gethin, Alix Myczkowski, and Anne-Sophie Robilliard (2022a). “How Unequal Is Africa? Towards Distributional National Accounts, 1990-2019”. World Inequality Lab Working Paper.
- Chancel, Lucas and Thomas Piketty (2019). “Indian Income Inequality, 1922-2015: From British Raj to Billionaire Raj?” In: *Review of Income and Wealth* 65.S1, S33–S62.
- (2021). “Global Income Inequality, 1820–2020: the Persistence and Mutation of Extreme Inequality”. In: *Journal of the European Economic Association* 19.6, 3025–3062.


- Chancel, Lucas, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman (2022b). *World Inequality Report 2022*. Tech. rep.
- Chen, Shaohua and Martin Ravallion (2010). “The Developing World Is Poorer Than We Thought, But No Less Successful In the Fight Against Poverty”. In: *Quarterly Journal of Economics* 125.4, pp. 1577–1625.
- Chong, Alberto, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer (2014). “Letter Grading Government Efficiency”. In: *Journal of the European Economic Association* 12.2, pp. 277–299.
- Cutler, David M., Kaushik Ghosh, Kassandra L. Messer, Trivellore Raghunathan, Allison B. Rosen, and Susan T. Stewart (2022). “A Satellite Account for Health in the United States”. In: *American Economic Review* 112.2, pp. 494–533.
- Das, Jishnu, Alaka Holla, Aakash Mohpal, and Karthik Muralidharan (2016). “Quality and Accountability in Health Care Delivery: Audit-Study Evidence from Primary Care in India”. In: *American Economic Review* 106.12, pp. 3765–3799.
- De Rosa, Mauricio, Ignacio Flores, and Marc Morgan (2022). “More Unequal or Not as Rich? Revisiting the Latin American Exception”. Stone Center for Inequality Studies Working Paper Series.
- Deaton, Angus (2005). “Measuring Poverty in a Growing World (Or Measuring Growth in a Poor World)”. In: *The Review of Economics and Statistics* 87.1, pp. 1–19.
- (2010). “Price Indexes, Inequality, and the Measurement of World Poverty”. In: *American Economic Review* 100.1, pp. 5–34.
- Eshaghnia, Sadegh, James J. Heckman, and Goya Razavi (2021). “The Willingness to Pay for School Quality, Neighborhood Attributes and Later Life Outcomes”. Working Paper.
- Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer (2020). “Welfare Analysis Meets Causal Inference”. In: *Journal of Economic Perspectives* 34.4, pp. 146–167.
- Fisher-Post, Matthew and Amory Gethin (2023). “Government Redistribution and Development: Global Estimates of Tax-and-Transfer Progressivity, 1980-2019”. Working Paper.
- Fowler, Anthony (2013). “Electoral and Policy Consequences of Voter Turnout: Evidence from Compulsory Voting in Australia”. In: *Quarterly Journal of Political Science* 8.1, pp. 159–182.
- Fujiwara, Thomas (2015). “Voting Technology, Political Responsiveness, and Infant Health: Evidence from Brazil”. In: *Econometrica* 83.2, pp. 423–464.
- Gadenne, Lucie, Samuel Norris, Monica Singhal, and Sandip Sukhtankar (2022). “In-Kind Transfers As Insurance”. Working Paper.
- GBD (2022). “Assessing performance of the Healthcare Access and Quality Index, overall and by select age groups, for 204 countries and territories, 1990–2019: a systematic analysis from the Global Burden of Disease Study 2019”. In: *Lancet Glob Health* 10.12, pp. 1715–1743.

- Germain, Jean-Marc, Lucas Chancel, Mathias André, and Thomas Blanchet (2021). *Rapport du groupe d'experts sur la mesure des inégalités et de la redistribution*. Tech. rep. INSEE.
- Gethin, Amory (2023a). “Distributional Growth Accounting: Education and the Reduction of Global Poverty, 1980-2022”. Working Paper.
- (2023b). “Who Benefits from Public Goods? Evidence from South Africa”. Working Paper.
- Gethin, Amory, Samuel Kofi Tetteh Baah, and Christoph Lakner (2023). “Global Schooling Inequality, 1980-2022”. Working Paper.
- Glassman, Brian (2019). “Multidimensional Deprivation in the United States: 2017”. American Community Survey Report.
- Goldman, Maya, Ingrid Woolard, and Jon Jellema (2020). “The Impact of Taxes and Transfers on Poverty and Income Distribution in South Africa 2014/2015”. AFD Research Paper.
- Hanushek, Eric A., Jens Ruhose, and Ludger Woessman (2017). “Knowledge Capital and Aggregate Income Differences: Development Accounting for US States”. In: *American Economic Journal: Macroeconomics* 9.4, pp. 184–224.
- Hendren, Nathaniel and Ben Sprung-Keyser (2020). “A Unified Welfare Analysis of Government Policies”. In: *The Quarterly Journal of Economics* 135.3, pp. 1209–1318.
- Herrera, Santiago and Abdoulaye Ouedraogo (2018). “Efficiency of Public Spending in Education, Health, and Infrastructure”. World Bank Policy Research Paper.
- Heston, Alan (2013). “Government Services: Productivity Adjustments”. In: *Measuring the Real Size of the World Economy: The Framework, Methodology, and Results of the International Comparison Program – ICP*. Ed. by World Bank. The World Bank. Chap. 11, pp. 301–318.
- Jestl, Stefan and Emanuel List (2022). “Inequality, Redistribution, and the Financial Crisis: Evidence from Distributional National Accounts for Austria”. In: *The Review of Income and Wealth* forthcoming.
- Khemani, Stuti, James Habyarimana, and Irfan Nooruddin (2019). “What Do Poor People Think About Direct Cash Transfers?” Brookings Future Development.
- Lakner, Christoph and Branko Milanovic (2016). “Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession”. In: *The World Bank Economic Review* 30.2, pp. 203–232.
- Lindert, Peter H. (1994). “The Rise of Social Spending, 1880-1930”. In: *Explorations in Economic History* 31.1, pp. 1–37.
- Liscow, Zachary and Abigail Pershing (2022). “Why Is So Much Redistribution In-Kind and Not in Cash? Evidence from a Survey Experiment”. In: *National Tax Journal* 75.2.
- Lobeck, Max and Morten Nyborg Støstad (2022). “The Consequences of Inequality: Beliefs and Redistributive Preferences”. Working Paper.

- Lokshin, Michael, Martin Ravallion, and Iván Torre (2022). “Is Social Protection a Luxury Good?” World Bank Policy Research Working Paper.
- Lustig, Nora, ed. (2018). *Commitment to Equity Handbook: Estimating the Impact of Fiscal Policy on Inequality and Poverty*. Brookings Institution Press.
- Mauro, Paolo, Rafael Romeu, Ariel Binder, and Asad Zaman (2015). “A modern history of fiscal prudence and profligacy”. In: *Journal of Monetary Economics* 76.1, pp. 55–70.
- Meltzer, Allan H. and Scott F. Richard (1981). “A Rational Theory of the Size of Government”. In: *Journal of Political Economy* 89.5, pp. 914–927.
- Mladenka, Kenneth R. and Kim Quaille Hill (1978). “The Distribution of Urban Police Services”. In: *The Journal of Politics* 40.1, pp. 112–133.
- Morgan, Marc (2017). “Extreme and Persistent Inequality: New Evidence for Brazil Combining National Accounts, Surveys and Fiscal Data, 2001-2015”.
- Muralidharan, Karthik (2019). “The State and the Market in Education Provision: Evidence and the Way Ahead”. Columbia SIPA Working Paper 2019-06.
- Muralidharan, Karthik and Venkatesh Sundararaman (2015). “The Aggregate Effect of School Choice: Evidence from a Two-Stage Experiment in India”. In: *The Quarterly Journal of Economics* 130.3, 1011–1066.
- O’Dea, Cormac and Ian Preston (2010). “The distributional impact of public spending in the UK”. IFS Working Paper W12/06.
- Paulus, Alari, Holly Sutherland, and Panos Tsakloglou (2010). “The distributional impact of in-kind public benefits in European countries”. In: *Journal of Policy Analysis and Management* 29.2, pp. 243–266.
- Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). “Distributional National Accounts: Methods and Estimates for the United States”. In: *The Quarterly Journal of Economics* 133.2, pp. 553–609.
- Piketty, Thomas, Li Yang, and Gabriel Zucman (2019). “Capital Accumulation, Private Property, and Rising Inequality in China, 1978-2015”. In: *American Economic Review* 109.7, pp. 2469–2496.
- Pinkovskiy, Maxim and Xavier Sala-i-Martin (2016). “Lights, Camera ... Income! Illuminating the National Accounts-Household Surveys Debate”. In: *The Quarterly Journal of Economics* 131.2, pp. 579–631.
- Ravallion, Martin (2012). “Why Don’t We See Poverty Convergence?” In: *American Economic Review* 102.1, pp. 504–523.
- Riedel, Lukas and Stichnoth Holger (2022). “Government expenditure in the DINA framework: Allocation methods and consequences for post-tax income inequality”. ZEW Discussion Papers, No. 22-004.

- Sala-i-Martin, Xavier (2006). “The World Distribution of Income: Falling Poverty and... Convergence, Period”. In: *The Quarterly Journal of Economics* 121.2, pp. 351–397.
- Thesmar, David and Augustin Landier (2022). *Le prix de nos valeurs*. Flammarion.
- Venkat, Hamsa and Nic Spaul (2015). “What do we know about primary teachers’ mathematical content knowledge in South Africa? An analysis of SACMEQ 2007”. In: *International Journal of Educational Development* 41.1, pp. 121–130.
- Verbist, Gerlinde, Michael Förster, and Maria Vaalavuo (2012). “The Impact of Publicly Provided Services on the Distribution of Resources: Review of New Results and Methods”. OECD Working Paper.
- Wagstaff, Adam, Marcel Bilger, Leander R. Buisman, and Caryn Bredenkamp (2014). “Who Benefits from Government Health Spending and Why?” World Bank Policy Research Working Paper 7044.
- Walter, Torsten Figueiredo (2020). “Misallocation in the Public Sector? Cross-Country Evidence from Two Million Primary Schools”. Working Paper.
- World Bank (2013). *Measuring the Real Size of the World Economy: The Framework, Methodology, and Results of the International Comparison Program – ICP*. The World Bank.
- (2018). *Poverty and Shared Prosperity 2018: Piecing Together the Poverty Puzzle*. Tech. rep.
- Yu, Bingxin, Eduardo Magalhaes, and Samuel Benin (2015). “A Database User Manual for SPEED: Statistics on Public Expenditures for Economic Development”. International Food Policy Research Institute.


Figure 1 – Global Poverty and Public Goods: Global Poverty Headcount Ratio, 1980-2019



Notes. The figure plots the evolution of the poverty headcount ratio at \$2.15 per day (2017 PPP USD) in the world as a whole, for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. Spikes correspond to lower and upper scenarios on the distribution of transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure 2 – Public Goods and Poverty Measurement


(a) Public and Private Goods are Substitutes



Notes. Author's computations combining national budget data (public health expenditure) and World Bank estimates (healthcare-driven poverty). The figure plots the relationship across countries between public health spending, expressed as a share of national income, and healthcare-driven poverty, measured as the share of the population falling into poverty due to out-of-pocket health expenditure. In countries spending more on public healthcare, fewer households fall into poverty due to own spending on healthcare.

Figure 2 – Public Goods and Poverty Measurement


(b) Public Goods Matter for Non-Monetary Dimensions of Quality of Life



45


Notes. Author's computations combining national budget data (public expenditure), World Bank estimates (monetary poverty rate), and Oxford Poverty and Human Development Initiative estimates (multidimensional poverty rate). The figure plots the relationship across countries between public expenditure on education, health, housing, and community services, and the gap between monetary and multidimensional measures of poverty. Monetary poverty: share of population spending less than \$2.15 per day (2017 PPP USD). Multidimensional poverty: index combining deprivation in health, education, and living standards (see [Alkire, Kanagaratnam, and Suppa, 2021](#)). In countries with greater spending on basic public services, fewer households fall in multidimensional poverty relative to those falling in monetary poverty.

Figure 3 – Public Goods Are Progressive
Distribution of Share of Transfers Received by the Bottom 50%




Notes. The figure represents the distribution of the share of government transfers and the share of pretax income received by the bottom 50% of the pretax income distribution in each country.

Figure 4 – Public Goods Have Grown
Global Real Public Expenditure Per Capita, 1980-2019



Notes. Author's computations using national budget data. The figure represents the evolution of general government expenditure per capita by function, expressed in 2021 PPP US dollars, in the world as a whole. Economic affairs include recreation and culture.


Figure 4 – Public Goods Have Grown
Expenditure on Public Goods by Country Income Group, 1980-2019



Notes. Author's computations combining national budget data. The figure represents the average share of national income spent on social assistance and public goods by country income group. Population-weighted averages across all countries in each group. See appendix figure A.5.38 for the composition of country income groups.

Figure 5 – The Distribution of Public Goods in International Perspective


(a) Public Goods Received by the Bottom 50% Around the World



Notes. The figure maps total in-kind transfers received by the bottom 50% in each country in 2019, expressed as a share of national income. The unit of observation is the individual. Income is split equally between all household members.


Figure 5 – The Distribution of Public Goods in International Perspective

(b) Public Goods Received by the Bottom 50% in Selected Countries




Notes. The figure shows the level and composition of in-kind transfers received by the bottom 50% in each country or region in 2019, expressed as a share of national income. The unit of observation is the individual. Income is split equally between all household members.

Figure 6 – Public Goods Received by the Global Bottom 20%, 1980-2019 (% of Global Income)



Notes. The figure plots the level and composition of public goods accruing to the global bottom 20%, expressed as a share of global income. The unit of observation is the individual. Income is split equally between all household members.


Figure 7 – Public Goods and Global Poverty Reduction:
Real Average Income of the Global Bottom 20%, 1980-2019



Notes. The figure plots the evolution of the global bottom 20% real average income from 1980 to 2019, before and after accounting for cash transfers and public goods. The unit of observation is the individual. Income is split equally between all household members.

Figure 8 – Global Inequality and Public Goods

(a) Real Income Growth Rate by Global Income Percentile, 1980-2019




53

Notes. The figure plots total real income growth by global income percentile from 1980 to 2019 for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. Capped spikes correspond to lower and upper scenarios on the progressivity of transfers. The unit of observation is the individual. Income is split equally between all household members.


Figure 8 – Global Inequality and Public Goods

(b) Global Top 10% to Bottom 50% Average Income Ratio



Notes. The figure plots the ratio of the average income of the top 10% to that of the bottom 50% in the world as a whole, for different income concepts. Capped spikes correspond to lower and upper scenarios on the progressivity of transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure 9 – Decomposing the Incidence of Public Goods on Global Poverty



Notes. The figure plots the share of the world population living with less than \$2.15 per day in 2019, measured in 2017 PPP USD, by income concept. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The fourth bar assumes that all transfers are received on a lump sum basis: $\gamma(m_i) = \gamma$. The next bar further considers that all countries have welfare states similar to that of Nordic countries, that is, general government expenditure is set at 50% of national income in each country. The last bar considers that no taxes are paid to finance transfers. The unit of observation is the individual. Income is split equally between all household members.

Table 1 – Methodology Used to Distribute Global Government Expenditure

	Source / Method	Avg. Share of NNI (%)		Share of Transfer Received (%)		
		G^j		$(\gamma^j, \text{Bottom } 50\%)$		
		1980	2019	Min	Mean	Max
Social Assistance	WID/CEQ/ASPIRE	2.6%	2.9%	16%	64%	92%
Education	GKL	3.5%	4.4%	25%	46%	64%
Health	WID/CEQ	2.5%	3.5%	29%	50%	69%
All Others	Prop. disposable income	17.4%	17.8%	8%	16%	30%
Economic Affairs		6.3%	5.8%			
General Public Services		5.6%	5.5%			
Public Order & Safety		1.4%	2.0%			
Other		4.1%	4.6%			
Total		26.0%	28.6%	15%	29%	48%

Notes. The table reports the sources used to distribute global government expenditure, together with summary statistics on expenditure by function as a share of national income and the share of expenditure received by the bottom 50% in each country. GKL: [Gethin, Kofi Tetteh Baah, and Lakner \(2023\)](#). WID: [Blanchet, Chancel, and Gethin \(2022\)](#) for Europe and [Piketty, Saez, and Zucman \(2018\)](#) for the US. CEQ: Commitment to Equity Institute Database. Prop. disposable income: component distributed proportionally to posttax disposable income (pretax income, minus direct taxes, plus social assistance transfers).

Table 2 – Dimensions of Redistribution by Country Income Group and World Region

	Expenditure (% NNI) <i>G</i>	Share of Transfer Received (%) (γ , Bottom 50%)	Net Transfer Received (% NNI) (<i>g</i> , Bottom 50%)
Country Income Group			
Low-Income	23.3%	22.8%	5.3%
Lower-Middle-Income	26.3%	24.0%	6.3%
Upper-Middle-Income	25.6%	29.2%	7.4%
High-Income	30.4%	33.2%	10.1%
World Region			
Sub-Saharan Africa	25.9%	22.3%	5.7%
Middle East and Northern Africa	28.6%	25.5%	7.2%
China	23.3%	27.0%	6.3%
India	31.4%	19.2%	6.0%
Other Asia / Oceania	23.3%	27.7%	6.5%
Latin America	25.8%	30.1%	7.7%
US / Canada / Western Europe	30.3%	34.9%	10.6%

Notes. The table reports statistics on dimensions of in-kind redistribution by country income group (defined based on the World Bank's classification) and world region. All figures focus on public goods, that is, total government expenditure excluding social protection spending.

Table 3 – Public Goods and the Geography of Global Inequality

	Pretax National Income	Posttax Disposable Income	Posttax National Income
Theil Decomposition			
Theil Index	1.13	0.98	0.89
Between-Country Component	30%	33%	39%
Within-Country Component	70%	67%	61%
Share in Global Bottom 20%			
India	18%	21%	24%
China	11%	11%	8%
Pakistan	19%	24%	31%
Bangladesh	19%	20%	30%
Ethiopia	58%	66%	74%
Nigeria	23%	28%	34%
Other Asia / Oceania	17%	17%	17%
Other Sub-Saharan Africa	62%	65%	67%
Middle East and Northern Africa	19%	19%	17%
Latin America	17%	11%	6%
US / Canada / Western Europe	7%	2%	0%

Notes. The table reports a Theil decomposition of global inequality into a between-country and a within-country component, as well as the geographical composition of the global bottom 20% in 2019, for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The unit of observation is the individual. Income is split equally between all household members.

Appendices

A Additional Methodological Details	60
B Accounting for Public Sector Productivity	66
C Additional Figures and Tables	73

A. Additional Methodological Details

This section presents the methodology used to estimate the distribution of global pretax and posttax incomes. Section [A.1](#) outlines the data sources used. Section [A.2](#) explains the methodology used to construct aggregate government revenue and expenditure series. Section [A.3](#) covers the distribution of transfers.

A.1. Data Sources

A.1.1. Macroeconomic Aggregates

My main source for macroeconomic aggregates is the World Inequality Database (WID, see <http://wid.world>), which combines various data sources to provide harmonized national accounts series and population totals in all countries in the world from 1950 to 2021 ([Blanchet and Chancel, 2016](#)). I use five main variables from the WID database in my analysis: gross domestic products, net national incomes, total populations, national income deflators, and PPP conversion factors to 2021 US dollars.

A.1.2. Government Revenue Aggregates

For government revenue aggregates, I rely on [Bachas et al. \(2022\)](#), who build a new database on the level and composition of tax revenue in 150 countries since 1965. Their database provides information on total tax revenue as a share of net domestic product, together with a breakdown by type of tax (personal income taxes, corporate income taxes, social contributions, property and wealth taxes, indirect taxes, and other taxes).

A.1.3. Government Expenditure Aggregates

Estimating the evolution of consolidated government expenditure and its composition is challenging, and there exists no single data source providing harmonized information on spending on different policies across countries. Accordingly, I combine various data sources to build a new database on government expenditure by function.

My primary data source for total expenditure is [Mauro et al. \(2015\)](#), who draw on historical data from the IMF and other sources to construct a new database on total consolidated government expenditure as a share of GDP in 170 countries from 1800 to 2011 (59 countries are

covered in 1980, 91 in 1990, and 157 after 2000).³¹ The main advantage of this database is its historical coverage and conceptual consistency: total expenditure covers consolidated government, incorporating both central and local government expenditure. Its main limitation is that it does not provide any information on the composition of expenditure.

The main data source used to cover the composition of expenditure (as well as total expenditure after 2011) is the IMF, which provides data on spending by Classification of the Functions of Government (COFOG) in 172 countries. Depending on the country and year, the series cover either the general government, or only unconsolidated central, state, and local government expenditure.

I use other data sources on specific types of expenditure to complement and further decompose IMF data.

The IFPRI-SPEED database (Yu, Magalhaes, and Benin, 2015) covers total central government expenditure in 147 countries, incorporating some country-specific sources absent from IMF series.

The World Bank's World Development Indicators (WDI) database provides series on total education and health expenditure as a share of government spending in 208 countries.

For decomposing social protection expenditure into social insurance and social assistance, I rely on three sources: the OECD's Social Expenditure (SOCX) database, the United Nations Economic Commission for Latin America and the Caribbean's Social Expenditure database, and the World Bank's Atlas of Social Protection Indicators of Resilience and Equity (ASPIRE).³² All three datasets provide data on total social protection expenditure as a share of GDP, as well as its decomposition by type of program.

A.1.4. Pretax Income Distribution Data

Data on the distribution of pretax income by country since 1980 come from the World Inequality Database, which brings together country-specific studies (e.g., Piketty, Saez, and Zucman (2018) for the US and Blanchet, Chancel, and Gethin (2022) for Europe) and other data sources to provide estimates of average pretax income by generalized percentile in all countries around the world since 1980 (see Chancel and Piketty, 2021).

The income concept covered in pretax national income, that is, the sum of all personal income

³¹See <https://www.imf.org/external/datamapper/exp@FPP/USA>.

³²See <https://www.oecd.org/social/expenditure.htm>; https://statistics.cepal.org/portal/databank/index.html?lang=en&indicator_id=4407&area_id=; <https://www.worldbank.org/en/data/datatopics/aspire>.

flows before taking into account the operation of the tax-and-transfer system, but after taking into account the operation of pension and unemployment systems. By construction, average pretax income matches average net national income in each country.

A.1.5. Tax Incidence Data

For the distributional incidence of taxes, I rely on estimates from a companion paper ([Fisher-Post and Gethin, 2023](#)).

A.1.6. Transfer Incidence Data

For the distributional incidence of government expenditure, I rely on five data sources: [Piketty, Saez, and Zucman \(2018\)](#), [Blanchet, Chancel, and Gethin \(2022\)](#), [Gethin, Kofi Tetteh Baah, and Lakner \(2023\)](#), the CEQ database, and the World Bank's ASPIRE database.

[Piketty, Saez, and Zucman \(2018\)](#) provide in their microfile data on all cash and health transfers received by US individuals from 1962 to 2021. I use this information to compute the share of total cash and health transfers received by pretax income decile.

[Blanchet, Chancel, and Gethin \(2022\)](#) provide in their microfile data on family and social assistance transfers received by individuals in 32 European countries. I use it to compute the share of cash transfers received by pretax income decile in each country.

[Gethin, Kofi Tetteh Baah, and Lakner \(2023\)](#) provides unique information on school attendance by age and household income in 155 countries since 1980. I use it to compute the share of education expenditure received by income decile in each country.

The CEQ database provides estimates of the share of cash transfers, total education expenditure, and total health expenditure received by pretax income decile in 45 countries.

Finally, the World Bank's ASPIRE database draws on harmonized survey microdata to compute the share of social assistance transfers received by pretax income quintile in 108 countries over the 1998-2019 period (most countries are covered since the mid-2000s).

A.2. Harmonization of Government Expenditure by Function: G

I combine all available data sources to build a harmonized database on the level and composition of government expenditure since 1980. I proceed in two steps. First, I combine existing sources to estimate total consolidated government expenditure in all countries and years. Second, I estimate the composition of consolidated expenditure by function.

A.2.1. Total Government Expenditure

My primary data source to measure total consolidated government expenditure is [Mauro et al. \(2015\)](#), which I use for all country-years in which data is available. In countries not covered at all by [Mauro et al. \(2015\)](#), I use available IMF general government series. In countries not covered at all by any of these two sources, I use the sum of central, state, and local government expenditure reported in IMF series.

To cover all countries from 1980 to 2019, I then combine all data sources to carry these combined series backward and forward. First, I carry [Mauro et al. \(2015\)](#) series backward and forward using growth rates in IMF general government series as a share of GDP. When data is still missing, I use growth rates in IMF central, state, and local government. When data is still missing, I use growth rates in total tax revenue as a share of GDP from [Bachas et al. \(2022\)](#). When data is still missing, I use growth rates in central government expenditure as a share of GDP from the IFPRI-SPEED database. When data is still missing, I extrapolate total expenditure backwards and forwards as a constant share of GDP. Finally, in the 13 small countries with no data on total government expenditure at all, I take continental averages of total expenditure as a share of GDP.

A.2.2. Composition of Government Expenditure

As for total government expenditure, I combine available data sources to estimate the composition of expenditure by function. My primary data source is the IMF series, which decompose expenditure into 10 large COFOG categories: social protection, education, health, recreation and culture, housing and community amenities, environmental protection, economic affairs, public order and safety, defense, and general public services.

I give priority to general government expenditure series, and use the sum of central, state, and local government expenditure series only when general government data is not available at all in a given country. I then extrapolate the composition of expenditure backward and forward so as to cover the entire 1980-2019 period. For countries with no data on the composition of expenditure, I take continental averages.

World Bank education and health expenditure series tend to be more consistent and cover more countries and years, so I incorporate them directly into these estimates. To do so, I simply replace education and health expenditure as a share of the general government budget by World Bank series when available. I then proportionally adjust other components of the general government budget so that the share of expenditure going to each function sums up to 1. This ensures that the resulting education and health expenditure series are fully consistent with

World Bank data, while preserving the relative shares of other functions of government reported in IMF data.

Following the same principle, I then further decompose general public services and economic affairs into their subcomponents. As above, I use IMF series to split general public services into administration and debt service expenditure, extrapolating their respective ratios when data is missing. In countries with no data on these subcomponents, I assume that debt service absorbs one-third of general public services expenditure, which corresponds to the average observed across all country-years. I follow the same process to decompose economic affairs into transport expenditure and expenditure on other economic affairs.

Lastly, given that pretax income already includes pensions and unemployment benefits, I remove spending on social insurance transfers from social protection expenditure. To do so, I use the OECD's and the CEPAL's datasets to estimate a split between social insurance and social assistance transfers, and reduce social protection expenditure by the corresponding amount in the harmonized database. For countries not covered by these two datasets (all non-OECD, non-Latin American countries), I use the World Bank's ASPIRE database, which provides an estimate of total social assistance expenditure as a share of GDP in 124 countries. I take the ratio of this estimate to total social protection expenditure in my harmonized series, so as to reduce social protection expenditure to only cover social assistance. Finally, in countries with no data from either the OECD, the CEPAL or the World Bank, I make the conservative assumption that social protection expenditure matches social assistance expenditure (in other words, that the share of social insurance expenditure in social protection expenditure is zero).

A.3. Distribution of Transfers: $\gamma(m_i)$

I combine available data sources to estimate transfer incidence profiles by income group. My measure of interest consists in concentration curves, that is, the share of a specific type of transfer received by income decile.³³ I then distribute transfers by combining these profiles with government expenditure by function in each country.

In each case, I consider three scenarios for countries with missing data: one benchmark scenario corresponding to the average profile observed across all country-years; an upper bound in which missing countries are attributed the average transfer incidence profile of the five countries with the most progressive profiles; and a lower bound in which missing countries are attributed the average transfer incidence profile of the five countries with the most regressive profiles. In

³³Concentration curves are more meaningful to distribute transfers than incidence curves, given that unlike taxes paid, transfers received are not generally proportional to income or consumption.

the absence of consistent data on the evolution of transfer progressivity over time (with the exception of the United States), I assume that it has remained constant in each country.

Social Assistance I combine concentration curves of social assistance expenditure by pretax income decile or quintile from [Piketty, Saez, and Zucman \(2018\)](#), [Blanchet, Chancel, and Gethin \(2022\)](#), the World Bank's ASPIRE database, and the CEQ Institute, by order of priority. I then allocate total social assistance expenditure in each country-year based on these profiles.

Education For education, I derive concentration curves of education spending by combining data from [Gethin, Kofi Tetteh Baah, and Lakner \(2023\)](#) with series of public education expenditure per child by level from the UNESCO, as explained in section 3.3.

Health The CEQ database (and [Piketty, Saez, and Zucman, 2018](#) for health in the US) is, to the best of my knowledge, the only available data source providing consistent information on the distributional incidence of health expenditure. I allocate total health expenditure in each country-year based on the corresponding concentration curves by income decile.

Economic Affairs I assume that expenditure on economic affairs is received proportionally to consumption. I use incidence curves on the relationship between income and consumption from [Chancel et al. \(2022a\)](#), who combine a number of microdata sources to derive typical lower and upper bounds on savings rates by pretax income percentile. In my benchmark scenario, I apply the same consumption-income profile in each country, corresponding to the typical profile estimated in [Chancel et al. \(2022a\)](#). I then use their lower and upper bounds as lower and upper bounds on the progressivity of expenditure on economic affairs.

Other Government Expenditure Other components of the government budget include expenditure on economic affairs, public order and safety, housing and community amenities, administration, recreation and culture, defense, and environmental protection. I distribute them proportionally to posttax disposable income.

Distribution of Taxes I borrow estimates of the distribution of taxes by generalized percentile directly from [Fisher-Post and Gethin \(2023\)](#).

Debt Service, Budget Balance, and Local Taxes Finally, to reach a concept of posttax income consistent with the distributional national accounts framework ([Piketty, Saez, and Zucman,](#)

2018), I distribute debt service expenditure, the budget balance, and local taxes to individuals. This ensure that average income is consistent with the net national income. The main issue is that data on tax revenue from [Bachas et al. \(2022\)](#) only covers taxes collected by the central government. As a result, the gap between total consolidated government expenditure and central government revenue incorporates both local taxes and the government deficit, which available data do not allow to distinguish. In the absence of better information, I distribute the gap between total revenue and total expenditure proportionally to pretax income in each country.

B. Accounting for Public Sector Productivity

B.1. Conceptual Framework

I consider an extension in which the value of public goods is allowed to differ from cost of provision. The value of public goods received by individuals can theoretically be broken down into three components:

$$g(m_i) = \sum_j G^j \times \gamma^j(m_i) \times \theta^j(m_i) \quad (6)$$

With G^j government expenditure and $\gamma^j(m_i)$ the share of expenditure received by i . $\theta^j(m_i)$ captures the fact that for a given cost of provision, individuals may receive services of different quality. Empirically, it is useful to make a distinction between two notions of productivity:

$$\theta^j(m_i) = \Theta^j \times q^j(m_i) \quad (7)$$

Θ^j is the *aggregate productivity* of expenditure on function j , which does not depend on m_i . It captures the fact that the government may be more or less efficient at providing a given service than a benchmark production unit. For instance, public schools in country A may be on average less cost-efficient than public schools in country B, which implies that all public education transfers should be reduced by a constant factor in country A.

$q^j(m_i)$ is a *heterogeneous productivity* parameter. It captures the fact that the quality of services provided, holding cost constant, may differ between income groups. For instance, teachers teaching in poorer areas may be more or less qualified than those teaching in richer areas, independently from the wages they receive.

Consider for example a government providing free public education at a cost of $G^j = \$1000 \times N$,

with N the size of the population. Because of inequalities in access to public education, however, the poorest 20% only receive \$500 per capita of funding: $\gamma^j(m_i) = 0.1$. Furthermore, the government appears to be particularly inefficient at providing public education: it underperforms by 50% relative to what it could do if it was at the production possibility frontier, which implies that $\Theta^j = 0.5$. Finally, schools attended by children belonging to the bottom quintile appear to be 20% less efficient at providing education than the average school in the country: $q^j(m_i) = 0.8$. Combining the different parameters, we get: $g^j(Q_1) = \$500 \times 0.5 \times 0.8 = \200 .

B.2. Aggregate Productivity Θ^j

I start with the estimation of aggregate productivity Θ^j , corresponding to the overall efficiency of the government at providing public services.

B.2.1. Methodology

Following the existing literature measuring the productivity of governments by combining data on outcomes with data on government expenditure (e.g., [Adam, Delis, and Kammas, 2011](#); [Afonso, Schuknecht, and Tanzi, 2005](#); [Herrera and Ouedraogo, 2018](#)), I propose to estimate Θ^j by benchmarking the productivity of governments around the world to one another. If a government produces more output than any other for a given cost, then its efficiency is set to 1, and the productivity of other governments with comparable costs is estimated based on the outputs they deliver. The advantage of this approach is its simplicity and transparency: governments delivering better education and health outcomes are considered to be more productive.

I estimate simple models of public sector productivity based on international data covering government expenditure and outcomes. In broad strokes, I choose a function of government (e.g., health) and collect cross-country data on expenditure (public health spending per capita), other inputs (e.g., GDP per capita), and an outcome of interest (mortality). I then use data envelopment analysis to non-parametrically estimate the technical frontier, defined as the maximum output ever achieved in any country-year for a given level of expenditure and other inputs (e.g., [Herrera and Ouedraogo, 2018](#)). Finally, I use the estimated frontier to estimate Θ^j , based on the extent to which output could be improved without changing costs in a given country-year.³⁴ This yields measures of technical efficiency ranging from 0 to 1 for each country-year covered by the data.

³⁴I use the `teradial` command in Stata.

I apply this methodology to estimate the productivity of public education and public healthcare. For each of these two functions of government, I estimate two alternative production frontiers: one based on a single input and a single output, and one that incorporates additional inputs to account for the fact that, for instance, education outcomes might be higher because of higher GDP per capita rather than greater education spending. For other public goods, given the absence of high-quality data on service delivery, I take the average of public education and public healthcare measures. Finally, I interpolate between years and extrapolate backwards and forwards measures of productivity by function in each country, so as to cover the 1980-2019 period. For countries with no data at all, I take the global average observed in each year.

B.2.2. Productivity of Public Education

Inputs The first element required to estimate public education productivity is a measure of cost of provision. I take public education expenditure per child, expressed in 2021 PPP US dollars, estimated from the public spending database compiled in this paper. For the estimation of multiple-input efficiency, I add three auxiliary inputs to the model: the log of GDP per capita in 2021 PPP USD (available from the WID), the log of the adult literacy rate, and the log of the share of children enrolled in private schools (both available from the World Bank's WDI)

Output The second element needed is a measure of government performance. Following the large literature in macroeconomics investigating the role of education in explaining differences in economic development (e.g., [Hanushek, Ruhose, and Woessman, 2017](#)), I propose to measure the output of the education system as the expected human capital that a child can hope to obtain at age 5:

$$Y^{\text{education}} = \exp(r_S S + r_Q Q) \quad (8)$$

With S expected years of schooling at age 5, r_S the return to a year of schooling, Q a measure of education quality, and r_Q the return to education quality. Data on expected years of schooling come from the UNESCO and covers 202 countries over the 1970-2020 period. Education quality is taken from [Altinok, Angrist, and Patrinos \(2018\)](#), who compile data from various international test scores to construct a new database of education quality in 134 countries. The return to schooling is set to 10% per year and the return to quality to 15% per standard deviation, following the existing literature.

Results Figure A.3.32 plots the resulting relationship between performance and cost of provision for all country-years. There is a very strong correlation between the two variables ($\rho = 0.9$, $R^2 = 0.82$): countries spending more on education display education systems of substantially better quality. Yet, there is also significant dispersion in the expected human capital stock achieved for a given level of government expenditure. The upper dashed line represents the efficient frontier, estimated using data envelopment analysis with variable returns to scale. This corresponds to a piecewise linear estimate of the maximum achievable output by level of expenditure.

The trajectories of Niger, Indonesia, and South Korea are represented as examples. Education expenditure and schooling outcomes have significantly increased during this period in all three countries. Niger stands quite far below the frontier, while South Korea has remained one of the most cost-efficient countries in the database throughout the period. Indonesia falls somewhat in-between. The corresponding measures of education productivity are about $\Theta^{\text{education}} = 0.5$ for Niger, 0.85 for Indonesia, and 0.9 for South Korea in the last year available.³⁵

B.2.3. Productivity of Public Healthcare

Inputs As for public education, the first step is to collect data on cost of provision. Given the particular role that private healthcare can play in some countries, I focus on total healthcare expenditure per capita (private and public combined). For the estimation of multiple-input efficiency, I add two auxiliary inputs to the model: the log of GDP per capita in 2021 PPP USD and the share of private health expenditure in total current health expenditure (both available from the WDI). All data series come from the World Bank’s World Development Indicators.

Output Finding an accurate measure of the quality of healthcare provision is more challenging than for education. Indeed, unlike the human capital stock, which has a clear cardinal (monetary) interpretation, there is no obvious measure of healthcare performance whose units are directly comparable to cost of provision. Quality-adjusted life expectancy is often taken as a measure of interest (e.g., Cutler et al., 2022), yet this indicator is, by itself, arguably a poor measure of the performance of the healthcare system. Given these limitations, I turn instead to the healthcare access and quality (HAQ) index estimated in the context of the global

³⁵Notice that as shown in figure A.3.32, I fit the efficiency frontier using the log of the human capital stock. To get correct efficiency measures, one then needs to convert the ratio of logs into the ratio of actual human capital stocks. More precisely, we have a measure of efficiency θ^{\log} such that: $\theta^{\log} = \frac{\log x}{\log \bar{f}(x)}$, with \bar{x} the technical frontier evaluated at x . The objective is to convert θ^{\log} into $\theta = \frac{x}{\bar{f}(x)}$. Rearranging yields $\bar{f}(x) = \exp(\frac{\log x}{\theta^{\log}})$ and hence $\theta = \frac{x}{\exp(\frac{\log x}{\theta^{\log}})}$.

burden of disease study (GBD, 2022). This indicator ranks healthcare systems from 0 to 100, based on death rates from 32 causes of death that could be avoided by timely and effective medical care. The main advantage of the HAQ index is that it was specifically created by health experts to measure the ability of healthcare systems to cure preventable diseases: it is explicitly a measure of performance. It also has the advantage of covering nearly all countries in the world since 1990. The disadvantage is that it is normalized from 0 to 100, so it has no cardinal interpretation. In the absence of better solution, I re-express the HAQ index in units of life expectancy by first regressing it on life expectancy at birth, and then normalizing it using the coefficient obtained. Reassuringly, this correction only marginally affects efficiency scores.³⁶

Results Figure A.3.33 plots the resulting relationship between healthcare performance and cost of provision for all country-years. As for education, there is a very strong correlation between the two variables ($\rho = 0.93$, $R^2 = 0.87$): countries spending more on healthcare are much more able to limit deaths from curable diseases. The upper dashed line represents the efficient frontier, while the trajectories of Sweden, China, and India are represented for the sake of illustration. India is significantly below the frontier (with an implied Θ^{health} below 0.6 in all years), while China and Sweden have remained among the best-performing countries throughout the period.

B.2.4. Discussion: Estimates of Θ^j as Lower Bounds on Government Productivity

I view these estimates as providing a *lower bound* on government productivity, especially in poor countries, for three main reasons.

First, national income purchasing power parity conversion factors do already account for government productivity (World Bank, 2013). Indeed, public sector productivity is adjusted for all government services in the Asia-Pacific, Western Asia, and Africa regions, using a Cobb-Douglas function that assumes that government employees are less productive in poor countries because of a lower and less efficient stock of capital equipment (Heston, 2013). In OECD countries and the European Union, further adjustments are made for health and education, combining indicators on the quantity and quality of services provided (Blades, 2013). Hence, the correction made here to account for aggregate productivity implies adjusting transfers downwards twice, once when using PPP conversion factors to correct for price differences across countries, and once when multiplying transfers received by Θ^j .

³⁶More specifically, I run a linear regression of life expectancy on the HAQ index, controlling for the log of GDP per capita, years of schooling of the working-age population, and country fixed effects. I then multiply the HAQ index by the coefficient obtained, so as to re-express it in “units of life expectancy.”

Second, the frontier approach implies by construction that Θ^j cannot be greater than 1, given that the maximum input-output combination ever observed in any country-year is given a score of 1. As a result, governments are assumed to never be more productive than the private sector for any kind of service provided ($\Theta^j = 1$ corresponds to a government exactly as cost efficient as the private sector).

Third, omitted variable bias is likely to drive estimates of Θ^j in poor countries significantly *downwards*. Indeed, poor countries are likely to have worse outcomes for a given level of government expenditure not only because of inefficiencies, but also because of a number of other confounding factors. These include lower incomes, greater inequality, more extreme weather conditions, or lower basic knowledge, which directly affect education and health outcomes independently from government investment. For all these reasons, overall government expenditure is likely to be more efficient in these countries than what the model suggests.

B.2.5. Validation: Correlates of Government Efficiency

Finally, a useful way of checking the reliability of my measures of government productivity is to compare them to existing indicators. Appendix table [A.3.7](#) shows that education and healthcare productivity are positively correlated with a number of indicators of government efficiency available from international sources and the literature. This is especially true of healthcare productivity, which is positively associated with a composite index of government effectiveness ($\rho = 0.57$ for single-input estimates), lower corruption ($\rho = 0.43$), and more transparent policy-making ($\rho = 0.34$). I also find a positive correlation between my measures of healthcare efficiency and the index of public sector productivity of [Chong et al. \(2014\)](#) ($\rho = 0.29$), who mail letters to 159 countries and argue that the rate of return of these letters to their original sender provides a simple and transparent measure of government productivity.

All four of my measures of productivity are also highly correlated with one another. In particular, the cross-country correlation between single-input and multiple-input estimates is 0.94 for education and 0.97 for healthcare.³⁷ In other words, accounting for other factors affecting the relationship between government expenditure and outcomes does not appear to significantly alter rankings of which countries are more or less efficient. I view these results as additional reassuring evidence that my estimates capture broad differences in government productivity across countries relatively well.

³⁷See appendix table [A.3.8](#), which provides raw pairwise correlations between measures.

B.3. Heterogeneous Productivity $q^j(m_i)$

Heterogeneity in productivity refers to the fact that the quality of public goods provided may vary by income group independently from their cost of provision, because, for instance, poorer geographical areas in a given country may provide public services in a more or less cost efficient way. Estimating heterogeneous productivity at a global scale is extraordinarily challenging, given the lack of high-quality data on service delivery by income group. In the absence of better information, I investigate using subjective perceptions of public services from international survey data to derive estimates of heterogeneous productivity by income group around the world. The data source is the Gallup World Poll, a yearly survey conducted since 2005 in 165 countries, which asks respondents whether they are satisfied with different types of public services in their area. I aggregate average responses by income quintile to measure differences in satisfaction with local public education, healthcare, police, and transport services.³⁸ I then use relative responses as a scaling parameter, to increase or decrease the transfer received by each income group, for each of these four functions of government.

These subjective indicators have advantages and disadvantages. On the one hand, they are available for nearly all countries in the world and cover different types of public services, providing a simple and transparent measure of differences in the perceived quality of public services. On the other hand, they may suffer from significant measurement biases, in particular the fact that subjective perceptions may not be comparable across income groups because of differences in expectations of what “good” and “bad” public services might be. This could lead to underestimating inequalities in the quality of services received by income group, if richer respondents evaluate the quality of public services by comparing them to a higher benchmark than low-income households.


At the same time, existing studies suggest that heterogeneity in quality by income group remains relatively limited. Drawing on various data sources in the context of South Africa, [Gethin \(2023b\)](#) finds that inequalities in the quality of public services received by income group tend to be small, both for subjective or objective indicators. Subjective perceptions of public services also appear to track objective indicators of inequality in service delivery relatively well. Similarly, [Walter \(2020\)](#) provides evidence that pupil-teacher ratios tend to vary substantially within countries, in particular in developing countries, but that differences in local economic development or remoteness only explain a very small fraction of these variations.

³⁸Respondents are asked whether they are “Satisfied” or “Dissatisfied” with the public transportation system, the quality of roads and highways, the educational system or the schools, and the availability of quality health care. I use these four measures to derive estimates of heterogeneous productivity in the provision of transport, education, and health care. For police services, I rely on a question that asks whether respondents have “confidence in the local police force.”

C. Additional Figures and Tables


C.1. Additional Key Results

Figure A.1.1 – Validation of Methodology
In-Kind Transfers Received by the Bottom 50% in South Africa, Simplified versus Detailed Series




Notes. The figure plots the level and composition of public services received by the bottom 50% in South Africa, comparing simplified series (this paper) to detailed series constructed in [Gethin \(2023b\)](#).

Figure A.1.2 – Validation of Methodology
 Bottom 50% Share of Education Spending, Own Estimates Versus CEQ Database




Notes. The figure compares the share of education spending received by the bottom 50% for selected countries in the Commitment to Equity (CEQ) Database and in the database of [Gethin, Kofi Tetteh Baah, and Lakner \(2023\)](#).

Figure A.1.3 – Global Poverty Headcount Ratio at \$3.65 per day, 1980-2019




Notes. The figure plots the evolution of the poverty headcount ratio at \$3.65 per day (2017 PPP USD) in the world as a whole, for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.4 – Global Poverty Headcount Ratio at \$6.85 per day, 1980-2019




Notes. The figure plots the evolution of the poverty headcount ratio at \$6.85 per day (2017 PPP USD) in the world as a whole, for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.5 – Global Poverty Headcount Ratio at \$2.15 per day, 1980-2019 (World Bank Data)




Notes. The figure plots the evolution of the poverty headcount ratio at \$2.15 per day (2017 PPP USD) in the world as a whole, for different income concepts. Distributions of consumption or disposable income per capita from the World Bank. Pretax income is reconstructed as consumption minus social assistance plus direct taxes.

Figure A.1.6 – Global Poverty Headcount Ratio at \$3.65 per day, 1980-2019 (World Bank Data)




Notes. The figure plots the evolution of the poverty headcount ratio at \$3.65 per day (2017 PPP USD) in the world as a whole, for different income concepts. Distributions of consumption or disposable income per capita from the World Bank. Pretax income is reconstructed as consumption minus social assistance plus direct taxes.

Figure A.1.7 – Global Poverty Headcount Ratio at \$6.85 per day, 1980-2019 (World Bank Data)




Notes. The figure plots the evolution of the poverty headcount ratio at \$6.85 per day (2017 PPP USD) in the world as a whole, for different income concepts. Distributions of consumption or disposable income per capita from the World Bank. Pretax income is reconstructed as consumption minus social assistance plus direct taxes.

Figure A.1.8 – Real Average Income of the Global Bottom 50%, 1980-2019




Notes. The figure plots the evolution of the global bottom 50% real average income from 1980 to 2019, before and after accounting for cash transfers and public goods. The unit of observation is the individual. Income is split equally among all household members.

Figure A.1.9 – Total Growth in Posttax Disposable Income and Public Goods Received by Global Income Percentile, 1980-2019




Notes. The figure plots the total growth rate in real posttax disposable income and in the real value of public goods received by global posttax disposable income percentile from 1980 to 2019. The unit of observation is the individual. Income is split equally between all household members. A LOWESS smoothing with 0.5 bandwidth is applied to the public goods curve.

Figure A.1.10 – Total Expenditure on Public Goods Received by Global Income Decile, 1980-2019



Notes. The figure plots the share of global expenditure on public goods received by global income decile. The unit of observation is the individual. Income is split equally between all household members.


Figure A.1.11 – Total Expenditure on Public Goods Received by the Global Bottom 50% and Top 5%, 1980-2019



88


Notes. The figure plots total expenditure on public goods received by the bottom 50% and top 5% of earners in the world as a whole, expressed as a share of global income. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.12 – Level and Composition of Public Services Received by the Global Bottom 50%, 1980-2019 (% of Global Income)




Notes. The figure plots the share of global income accruing to the global bottom 50% in the form of public goods. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.13 – Level and Composition of Public Services Received by the Global Bottom 20% (Real 2021 PPP USD), 1980-2019




Notes. The figure plots the evolution of public services accruing to the global bottom 20%, expressed in real 2021 PPP US dollars. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.14 – Level and Composition of Public Services Received by the Global Bottom 50% (Real 2021 PPP USD), 1980-2019




Notes. The figure plots the evolution of public services accruing to the global bottom 50%, expressed in real 2021 PPP US dollars. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.15 – Gini Index of Global Income Inequality, 1980-2019




Notes. The figure plots the evolution of the Gini index of global income inequality for different income concepts. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.16 – Theil Index of Global Income Inequality, 1980-2019




Notes. The figure plots the evolution of the Theil index of global income inequality for different income concepts. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.17 – Share of Global Income Inequality Explained by Between-Country Inequalities, 1980-2019



Notes. The figure plots the evolution of the share of global income inequality explained by differences in average incomes between countries, computed from a Theil decomposition of global inequality into a between-country component and a within-country component. The unit of observation is the individual. Income is split equally between all household members.


Figure A.1.18 – Public Goods Received by the Global Bottom 20%: With Productivity-Adjusted Estimates



06


Notes. The figure plots the level of public goods accruing to the global bottom 20%, expressed as a share of global income, before and after adjusting for aggregate and heterogeneous productivity. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.19 – Global Poverty Headcount Ratio, 1980-2019: With Productivity-Adjusted Estimates




Notes. The figure plots the evolution of the poverty headcount ratio at \$2.15 per day (2017 PPP USD) in the world as a whole, for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.20 – Real Income Growth Rate by Global Income Percentile, 1980-2019:
With Productivity-Adjusted Estimates




Notes. The figure plots total real income growth by global income percentile from 1980 to 2019 for different income concepts. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.21 – Decomposing the Incidence of Public Goods on Global Poverty:
With Productivity-Adjusted Estimates, \$2.15 Threshold




Notes. The figure plots the share of the world population living with less than \$2.15 per day in 2019, measured in 2017 PPP USD, by income concept. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The fourth bar assumes no heterogeneous productivity: $q^j(m_i) = 1$. The next bar further assumes no aggregate inefficiency: $\Theta^j = 1$. The next bar assumes that all transfers are received on a lump sum basis: $\gamma(m_i) = \gamma$. The next bar further considers that all countries have welfare states similar to that of Nordic countries, that is, general government expenditure is set at 50% of national income in each country. The last bar considers that no taxes are paid to finance transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.22 – Decomposing the Incidence of Public Goods on Global Poverty:
With Productivity-Adjusted Estimates, \$3.65 Threshold



Notes. The figure plots the share of the world population living with less than \$3.65 per day in 2019, measured in 2017 PPP USD, by income concept. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The fourth bar assumes no heterogeneous productivity: $q^j(m_i) = 1$. The next bar further assumes no aggregate inefficiency: $\Theta^j = 1$. The next bar assumes that all transfers are received on a lump sum basis: $\gamma(m_i) = \gamma$. The next bar further considers that all countries have welfare states similar to that of Nordic countries, that is, general government expenditure is set at 50% of national income in each country. The last bar considers that no taxes are paid to finance transfers. The unit of observation is the individual. Income is split equally between all household members.

Figure A.1.23 – Decomposing the Incidence of Public Goods on Global Poverty:
With Productivity-Adjusted Estimates, \$6.85 Threshold



Notes. The figure plots the share of the world population living with less than \$6.85 per day in 2019, measured in 2017 PPP USD, by income concept. Posttax disposable income removes direct taxes and adds cash transfers. Posttax national income removes all taxes and adds all cash and in-kind transfers. The fourth bar assumes no heterogeneous productivity: $q^j(m_i) = 1$. The next bar further assumes no aggregate inefficiency: $\Theta^j = 1$. The next bar assumes that all transfers are received on a lump sum basis: $\gamma(m_i) = \gamma$. The next bar further considers that all countries have welfare states similar to that of Nordic countries, that is, general government expenditure is set at 50% of national income in each country. The last bar considers that no taxes are paid to finance transfers. The unit of observation is the individual. Income is split equally between all household members.

Table A.1.1 – Pairwise Correlations Between Dimensions
of Government Redistribution Across Countries

	Cost	Progressivity	Aggregate Productivity	Heterogeneous Productivity	NNI per capita
Cost	1.00				
Progressivity	0.60***	1.00			
Aggregate Productivity	0.42***	0.59***	1.00		
Heterogeneous Productivity	0.08	0.49***	0.22***	1.00	
NNI per capita	0.56***	0.71***	0.63***	0.28***	1.00

Notes. The table reports raw correlation coefficients between different dimensions of government redistribution across countries. Cost (C^j) corresponds to total general government expenditure as a share of net national income. Progressivity ($\gamma^j(m_i)$) is measured as the share of total government expenditure received by the bottom 50% (excluding social security). Aggregate productivity (Θ^j) corresponds to single-input, output-oriented estimates for each function of government. Heterogeneous productivity is measured as the relative quality of public services received by the bottom 20% in each country. Statistics computed over all countries in the database ($N = 174$). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.1.2 – Public Goods and Global Poverty Reduction:
Sensitivity to Different Specifications and Geographical Restrictions

	Global Poverty Headcount Ratio at \$1.9 Per Day		
	1980	2019	2019-1980
All Countries			
Posttax Disposable Income	21.1%	10.7%	-50%
Posttax National Income: Benchmark	19.2%	7.1%	-63%
Posttax National Income: Only Education & Health	26.1%	9.6%	-63%
Posttax National Income: Other Public Goods Lump Sum	15.6%	3.3%	-79%
Excluding China			
Posttax Disposable Income	16.9%	9.4%	-44%
Posttax National Income	15.2%	7.1%	-53%
Excluding India			
Posttax Disposable Income	14.5%	8.7%	-40%
Posttax National Income	12.9%	5.7%	-56%
Excluding China & India			
Posttax Disposable Income	10.2%	7.4%	-27%
Posttax National Income	8.9%	5.7%	-36%

Notes. The table reports how results on the incidence of public goods on global poverty reduction vary depending on assumptions regarding the progressivity of public goods and geographical restrictions. Only Education and Health: only allocate education and health expenditure. Other Public Goods Lump Sum: allocate all public goods other than education and health on a lump sum basis.

Table A.1.3 – Public Goods Provision Over the Course of Development:
Before and After Adjusting for Productivity

	Expenditure (% NNI) <i>G</i>	Share of Transfer Received (%) (γ , Bottom 50%)	Net Transfer Received (% NNI) (<i>g</i> , Bottom 50%)	Adjusted for Productivity (<i>g</i> , Bottom 50%)
Country Income Group				
Low-Income	23.3%	21.0%	4.9%	3.0%
Lower-Middle-Income	26.3%	23.3%	6.1%	4.0%
Upper-Middle-Income	25.6%	28.1%	7.1%	5.2%
High-Income	30.4%	33.0%	10.0%	8.3%
World Region				
Sub-Saharan Africa	25.9%	20.9%	5.4%	3.2%
Middle East and Northern Africa	28.6%	24.7%	7.0%	5.1%
China	23.3%	25.4%	5.9%	5.0%
India	31.4%	18.6%	5.8%	3.4%
Other Asia / Oceania	23.3%	27.1%	6.4%	4.8%
Latin America	25.8%	28.3%	7.2%	5.1%
US / Canada / Western Europe	30.3%	35.0%	10.6%	8.9%

Notes. The table reports statistics on dimensions of in-kind redistribution by country income group (defined based on the World Bank's classification) and world region. All figures focus on public goods, that is, total government expenditure excluding social protection spending. The last column adjusts estimates for differences in aggregate and heterogeneous productivity across countries.

Table A.1.4 – Public Goods, Quality of Life, and the Gap Between Surveys and National Accounts

	Expected Years of Schooling		Youth Literacy		Secondary School Enrollment Rate		Infant Mortality		Life Expectancy	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: No FE										
GDP-Survey Gap	0.16*** (0.02)	0.06*** (0.01)	0.09*** (0.02)	0.02 (0.02)	0.29*** (0.03)	0.10*** (0.02)	0.57*** (0.07)	-0.01 (0.03)	0.06*** (0.01)	0.00 (0.00)
Educ./Health Spending		0.19*** (0.00)		0.11*** (0.01)		0.29*** (0.01)		0.60*** (0.01)		0.06*** (0.00)
Panel B: Country FE										
GDP-Survey Gap	-0.04*** (0.01)	-0.01 (0.02)	0.08*** (0.03)	0.04 (0.03)	0.13*** (0.03)	-0.10*** (0.03)	0.33*** (0.05)	0.02 (0.04)	0.03*** (0.01)	0.00 (0.01)
Educ./Health Spending		0.23*** (0.01)		0.04*** (0.01)		0.34*** (0.01)		0.67*** (0.01)		0.06*** (0.00)
N	1193	1194	285	285	1409	1409	1760	1760	1772	1772
Adj. R-squared	0.93	0.88	0.86	0.87	0.82	0.89	0.88	0.95	0.87	0.92

Notes. Each column presents coefficients of a regression of a selected dependent variable on the gap between GDP and survey means, before and after controlling for education or health spending. GDP-Survey Gap: percentage difference between GDP per capita and survey mean income. Educ./Health Spending: log of public education spending (expected years of schooling, youth literacy, secondary school enrollment rate) or log of public health spending (infant mortality, life expectancy). Panel A runs simple OLS regressions. Panel B includes country fixed effects.


Table A.1.5 – Political Correlates of Public Goods Redistribution

	(1)	(2)	(3)	(4)	(5)	(6)
Electoral Democracy Index (0-1)	1.212*** (0.280)	1.423*** (0.295)	0.975*** (0.331)	0.745** (0.317)	1.155*** (0.338)	0.986*** (0.350)
Political Competition Index (0-10)	-0.032* (0.018)	-0.042** (0.019)	0.011 (0.022)	-0.010 (0.016)	0.003 (0.017)	0.013 (0.017)
Public Sector Corruption Index (0-1)	-0.583** (0.230)	-0.412* (0.248)	-0.581** (0.266)	0.254 (0.284)	0.355 (0.310)	0.366 (0.319)
Government Effectiveness (0-1)	-1.116*** (0.395)	-0.761* (0.424)	-1.816*** (0.491)	-0.689 (0.470)	-0.601 (0.502)	0.413 (0.551)
Log GDP Per Capita	0.784*** (0.061)	0.732*** (0.065)	0.761*** (0.071)	0.283** (0.126)	0.159 (0.136)	0.092 (0.143)
Additional Controls	X	X	X	X	X	X
Country FE				X	X	X
Excl. Western Democracies			X			X
Sample	1980-2019	2000-2019	2000-2019	1980-2019	2000-2019	2000-2019
N	2915	2637	2089	2915	2637	2089
Adj. R-squared	0.65	0.64	0.48	0.90	0.91	0.88

Notes. The table reports the results of a linear regression of redistribution on a number of political and economic variables. Redistribution is measured as the share of national income received by the bottom 50% in the form of public services. All estimates include country and year fixed effects and control for the following additional variables: bottom 50% pretax income share, log of total population, share of population aged 0-19, 20-39, and 40-59, and trade to GDP ratio. Country FE: country fixed effects. Excl. Western Democracies: excludes Western European countries, Canada, the United States, New Zealand, and Australia from the sample.


C.2. Macroeconomic Aggregates

Figure A.2.24 – Data Coverage of Total Government Expenditure and Revenue by Source




Notes. The figure shows the share of the world population covered by the different sources used to construct harmonized general government expenditure and central government revenue (in the case of Bachas et al. 2022) series.

Figure A.2.25 – Global Government Expenditure, 1980-2019 (% of Global Income)




Notes. The figure shows the evolution of real average global general government expenditure, expressed as a share of total global national incomes.

Figure A.2.26 – Government Expenditure on Public Goods Per Capita by Country Income Group, 1980-2019




Notes. The figure shows the evolution of average real per capita general government expenditure on public goods by country income group, expressed in 2021 PPP USD. Population-weighted average across all countries in each group.

Figure A.2.27 – Government Expenditure on Public Goods by Country Income Group, 1980-2019 (% of NNI)




Notes. The figure shows the evolution of average general government expenditure on public goods by country income group, expressed as a share of national income. Population-weighted average across all countries in each group.

Figure A.2.28 – Government Expenditure on Public Goods Per Capita by World Region, 1980-2019




Notes. The figure shows the evolution of real per capita general government expenditure on public goods by world region, expressed in 2021 PPP USD. Other Western countries: United States, Canada, Australia, New Zealand. Population-weighted average across all countries in each group.

Figure A.2.29 – Government Expenditure on Public Goods by World Region, 1980-2019 (% of NNI)




Notes. The figure shows the evolution of general government expenditure on public goods by world region, expressed as a share of national income. Other Western countries: United States, Canada, Australia, New Zealand. Population-weighted average across all countries in each group.

Figure A.2.30 – Government Tax Revenue Per Capita by World Region, 1980-2019



Notes. The figure shows the evolution of real per capita central government tax revenue by world region, expressed in 2021 PPP USD. Other Western countries: United States, Canada, Australia, New Zealand.


Figure A.2.31 – Government Tax Revenue by World Region, 1980-2019 (% of Regional Income)



Notes. The figure shows the evolution of central government tax revenue by world region, expressed as a share of total regional income. Other Western countries: United States, Canada, Australia, New Zealand.


C.3. Public Sector Productivity: Aggregate Productivity

Figure A.3.32 – Education Expenditure and Expected Human Capital at Age 5




Notes. The unit of observation is the country-year. Data on expected years of schooling from the UNESCO. Data on education expenditure per child come from estimates presented in this paper.

Figure A.3.33 – Health Expenditure and Quality of Healthcare




Notes. The unit of observation is the country-year. Data on healthcare access and quality index from the Global Burden of Disease Study. Data on health expenditure from the World Bank.

Figure A.3.34 – Distribution of Aggregate Public Sector Productivity: Education



Notes. The figure plots the distribution of aggregate public sector productivity Θ^j for education expenditure, plotted across all country-years in the database, for each of the four models considered.

Figure A.3.35 – Distribution of Aggregate Public Sector Productivity: Health



Notes. The figure plots the distribution of aggregate public sector productivity Θ^j for health expenditure, plotted across all country-years in the database, for each of the four models considered.

Table A.3.6 – Summary Statistics on Cross-Country Government Aggregate Productivity Measures

	Mean	SD	Min	Max
Education				
One Input	0.68	0.09	0.32	1.00
Multiple Inputs	0.73	0.09	0.39	1.00
Health				
One Input	0.71	0.19	0.12	1.00
Multiple Inputs	0.74	0.19	0.16	1.00

Notes. Statistics computed over all country-years in the database and weighted by total population.

Table A.3.7 – Correlates of Aggregate Government Productivity

	Education Single Input	Education Multiple Inputs	Health Single Input	Health Multiple Inputs	N
Chong et al. (2014) Mail Efficiency	0.08	0.00	0.29***	0.24***	159
Government Effectiveness	0.30***	0.13*	0.57***	0.49***	177
Control of Corruption	0.17**	0.04	0.43***	0.38***	177
Absence of Corruption	0.07	-0.05	0.27***	0.23***	160
Wastefulness of Government Spending	0.22***	0.14*	0.26***	0.24***	149
Irregular Payments and Bribes	0.24***	0.10	0.46***	0.41***	150
Favoritism in Government Decisions	0.15*	0.03	0.28***	0.22***	151
Transparency of Policymaking	0.20**	0.06	0.34***	0.29***	150
GDP per capita	0.35***	0.22***	0.62***	0.57***	177
Inequality in Public Service Delivery	0.34***	0.35***	0.30***	0.32***	160

Notes. The table reports raw pairwise correlations between the four measures of total technical efficiency and other qualitative indicators of government productivity. Correlations are computed over all countries with available data for each pair of indicators, for the last year available, and weighted by each country's total population. Chong et al. (2014) efficiency corresponds to the average number of days to get the letter back. GDP per capita data come from the World Inequality Database. Inequality in public service delivery is measured as the quality of public services received by the bottom quintile relative to the overall population ($q^j(Q1)$), estimated from the Gallup World Poll over the 2009-2021 period. Data on other indicators come from the World Bank. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


Table A.3.8 – Correlations Between Measures of Government Productivity

	Educ1	Educ2	Heal1	Heal2
Educ1		0.94***	0.57***	0.60***
Educ2	0.94***		0.54***	0.59***
Heal1	0.57***	0.54***		0.97***
Heal2	0.60***	0.59***	0.97***	

Notes. Correlations are computed over all countries with available data for each pair of indicators, for the last year available, and weighted by each country's total population. Educ: education; Heal: health. Numbers correspond to models with single (1) or multiple (2) inputs.

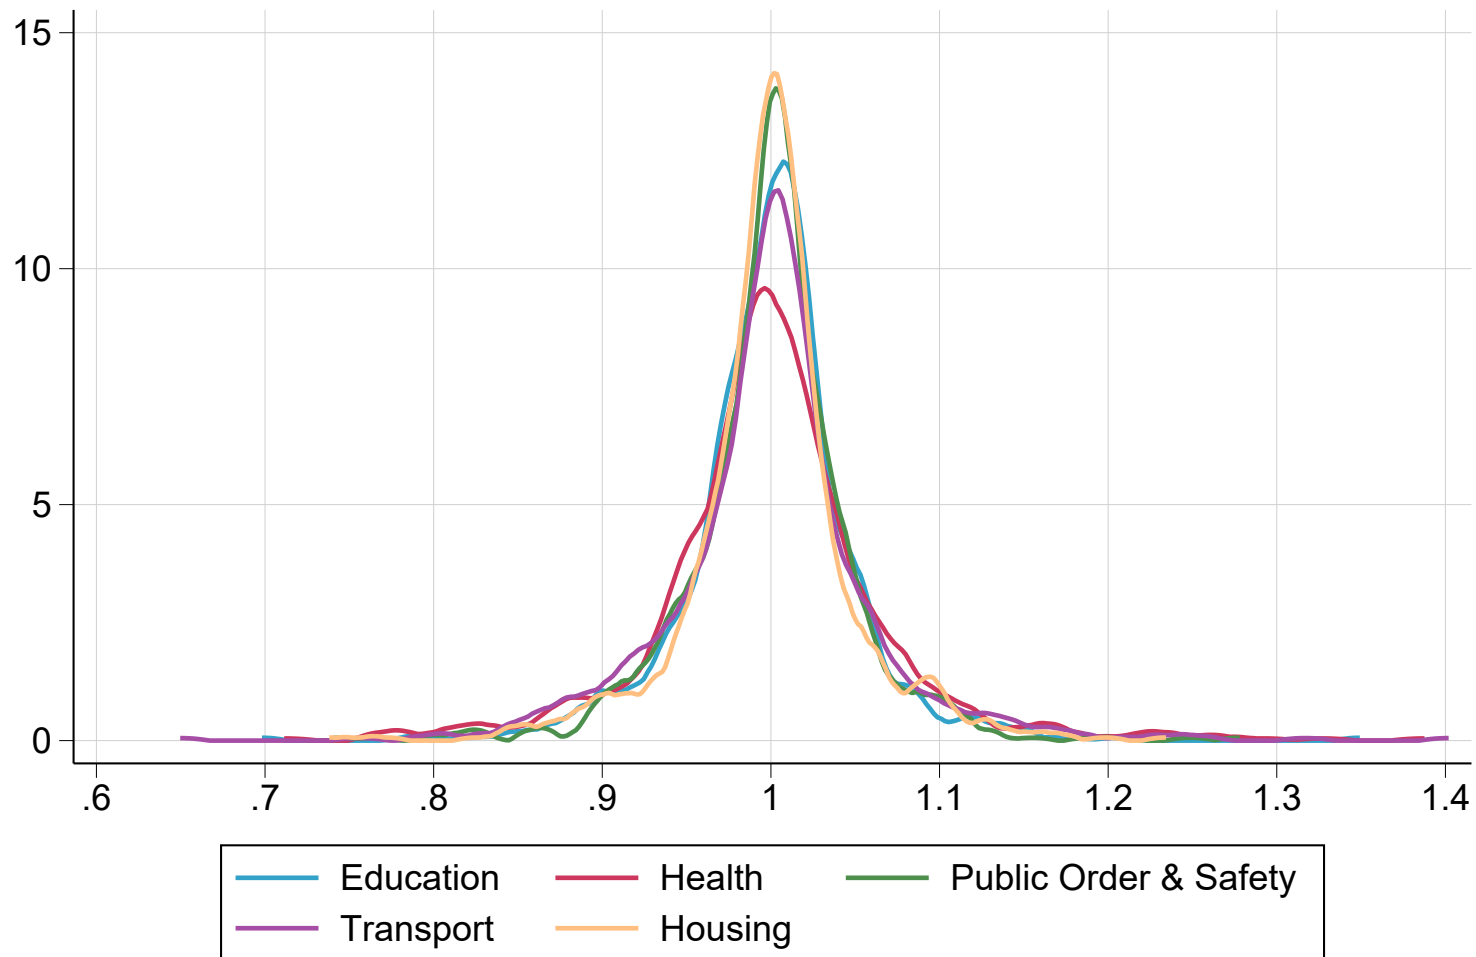
C.4. Public Sector Productivity: Heterogeneous Productivity

Figure A.4.36 – Average Heterogeneous Productivity Profiles by Function, World



Notes. Author's computations using Gallup World Poll data. The figure represents the average of heterogeneous productivity profiles $q^j(m_i)$ applied to correct in-kind transfers received by income quintile, computed over all countries over the entire 2009-2021 period. Numbers correspond to the ratio of the quality of the transfer received to average quality. Quality is measured as the share of respondents who declare being satisfied with public services in the city or area where they live, for the following services: public transportation systems, roads and highways, the educational system or the schools, the quality of water, and the availability of quality health care. The quality of police services is measured as the share of respondents who declare having confidence in the local police force.

Figure A.4.37 – Distribution of Heterogeneous Productivity Scores by Function



Notes. The figure represents the distribution of heterogeneous productivity scores by function, estimated from the Gallup World Poll data, across all countries with available data, for the bottom 20%. Figures correspond to the ratio of the quality of the transfer received by the bottom 20% to average quality. Quality is measured as the share of respondents who declare being satisfied with public services in the city or area where they live, for the following services: public transportation systems, roads and highways, the educational system or the schools, the quality of water, and the availability of quality health care. The quality of police services is measured as the share of respondents who declare having confidence in the local police force.


Table A.4.9 – Indicators of Heterogeneous Public Service Delivery by Income Quintile in South Africa

	Q1	Q2	Q3	Q4	Q5	$q^j(Q_1)$	Source
Subjective Indicators (% Positively Rating)							
Local public school	69%	69%	69%	68%	69%	1.01***	Census
Local public clinic	46%	45%	46%	46%	50%	0.98***	Census
Local public hospital	47%	47%	47%	48%	51%	0.97***	Census
Local police services	43%	43%	44%	45%	48%	0.97***	Census
Electricity supply	63%	63%	63%	64%	67%	0.99***	Census
Water supply	50%	54%	58%	62%	68%	0.85***	Census
Refuse removal services	49%	54%	57%	60%	66%	0.85***	Census
Sanitation services	52%	56%	59%	64%	74%	0.85***	Census
Government-subsidized dwelling	48%	49%	50%	51%	53%	0.96***	Census
Police response to reported crime	52%	53%	52%	53%	56%	0.98	VCS
Objective Indicators							
School teacher mathematics test success rate	38%	40%	40%	47%	67%	0.82***	SACMEQ
Share of reported crimes leading to arrest	24%	20%	21%	18%	20%	1.15	VCS
Asked to pay a bribe in past 12 months	5%	9%	8%	11%	15%	1.78***	VCS
Water interruption in past 3 months	19%	19%	17%	16%	14%	0.90***	Census
Electricity interruption in past 3 months	32%	28%	25%	21%	16%	0.76***	Census
Value of subsidized dwelling (R 1,000)	177	178	267	308	305	0.72***	GHS
Distance to Nearest Public Services (km)							
Primary school	1.5	1.5	1.6	1.8	2.0	1.12***	LCS
Secondary school	2.9	2.8	2.6	2.4	2.8	0.93***	LCS
Clinic	4.7	4.5	3.8	3.5	3.8	0.86***	LCS
Hospital	13.2	12.6	10.2	8.6	7.3	0.79***	LCS
Police station	8.6	8.1	6.1	4.9	4.6	0.75***	LCS
Public transport	1.1	1.0	1.1	1.0	1.3	1.04*	LCS

Notes. The table reports estimates of heterogeneous government productivity by income group, based on a number of subjective and objective indicators of public service delivery. Q1 to Q5 refer to income quintiles. $q^j(Q_1)$ is the corresponding measure of the relative quality of services received by the bottom quintile, equal to the ratio of the value of the indicator for Q1 to the overall sample mean (or its inverse when the scale of the variable is inverted). Statistical significance stars correspond to a regression of the indicator of interest on a dummy taking one if the individual belongs to the bottom quintile. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Census: 2016 national census. GHS: 2019 General Household Survey. VCS: 2017 Victims of Crime Survey. LCS: 2014-2015 Living Conditions Survey. SACMEQ: The Southern and Eastern Africa Consortium for Monitoring Educational Quality (estimates from [Venkat and Spaul, 2015](#)).


C.5. Maps

Figure A.5.38 – Country Income Groups




Notes. Authors' elaboration based on World Bank classification of country income groups.

Figure A.5.39 – General Government Expenditure, 2019 (% of National Income)



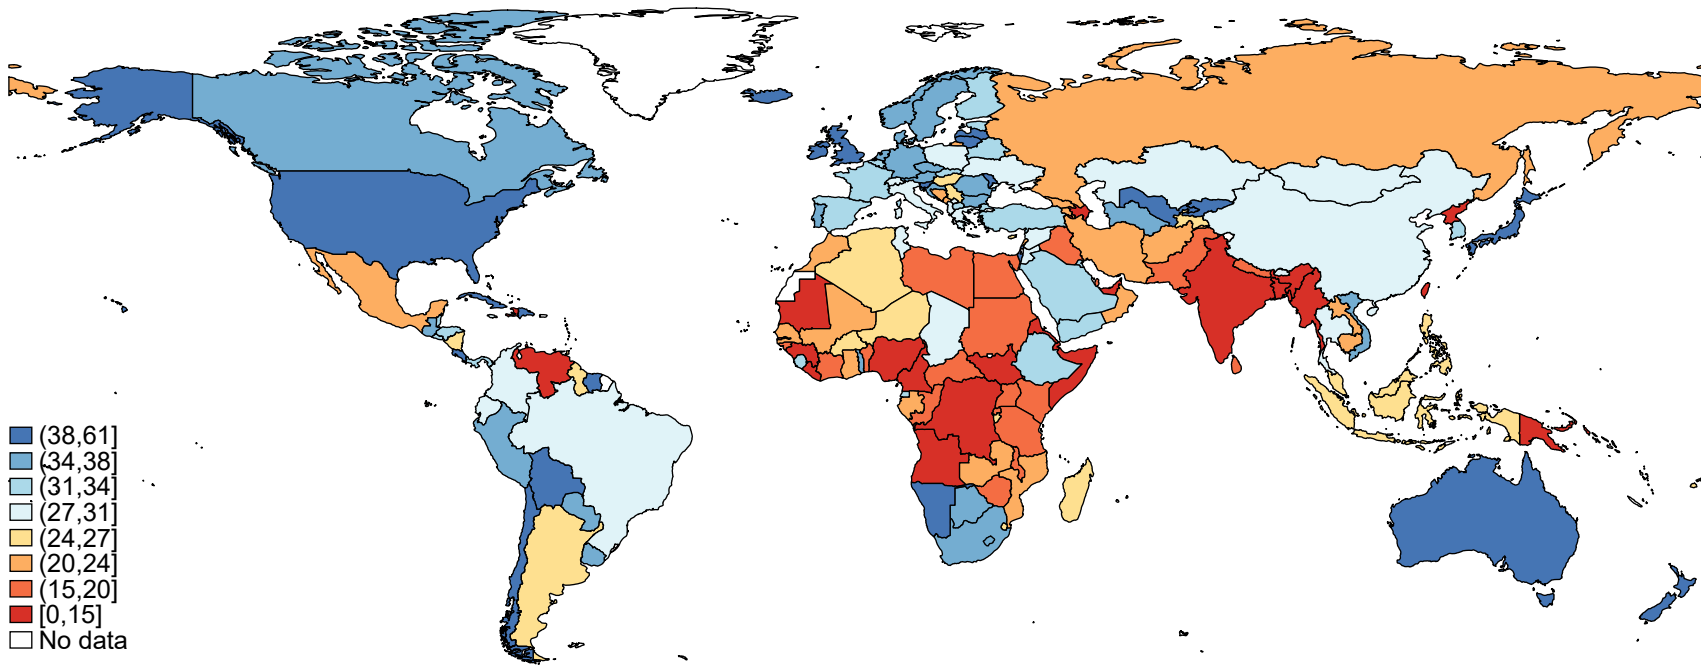
Notes. Authors' computations using national budget data.

Figure A.5.40 – Change in General Government Expenditure, 1980-2019 (% of National Income)



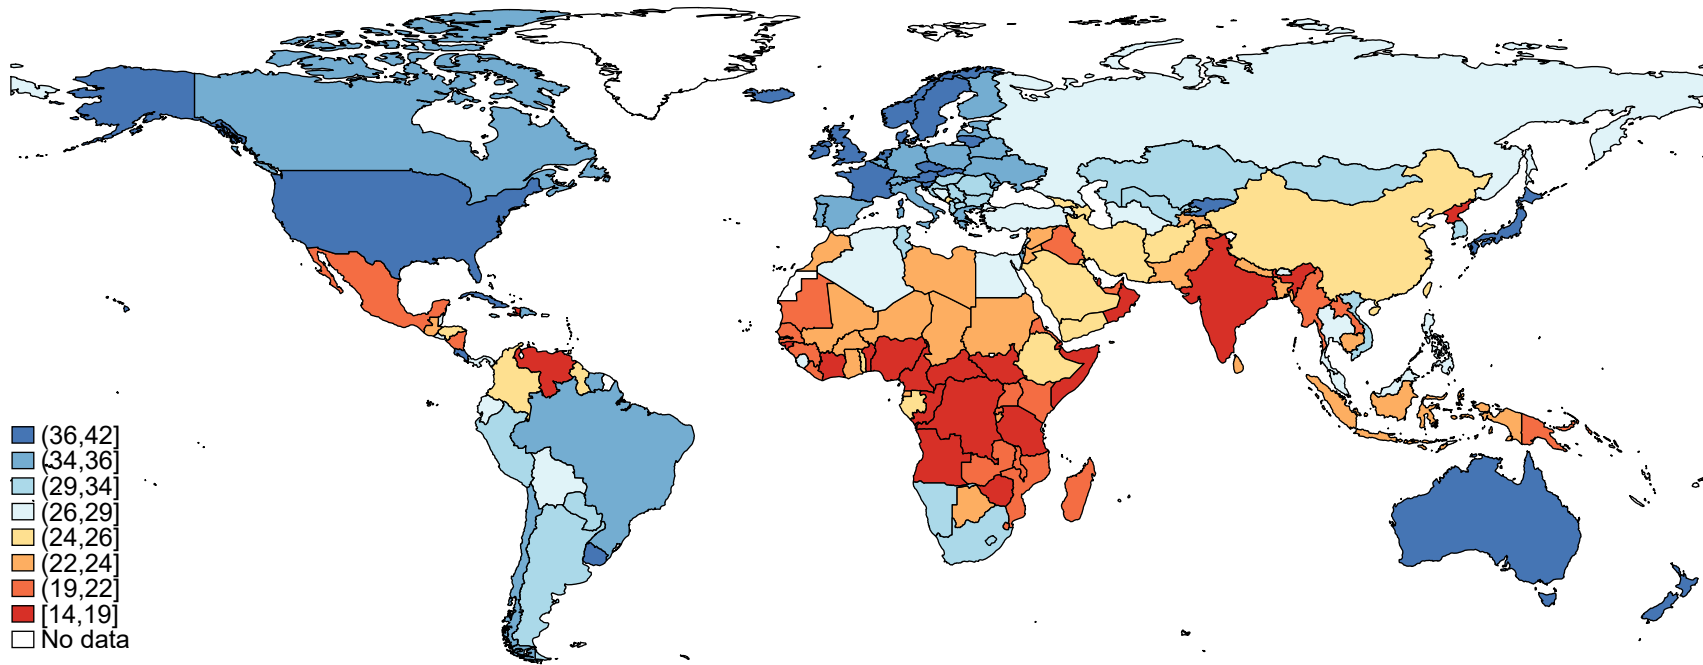
Notes. Authors' computations using national budget data.

Figure A.5.41 – General Government Expenditure on Education and Health, 2019 (% of Total Expenditure)




Notes. Authors' computations using national budget data.

Figure A.5.42 – Share of Expenditure on Public Goods Received by the Bottom 50%



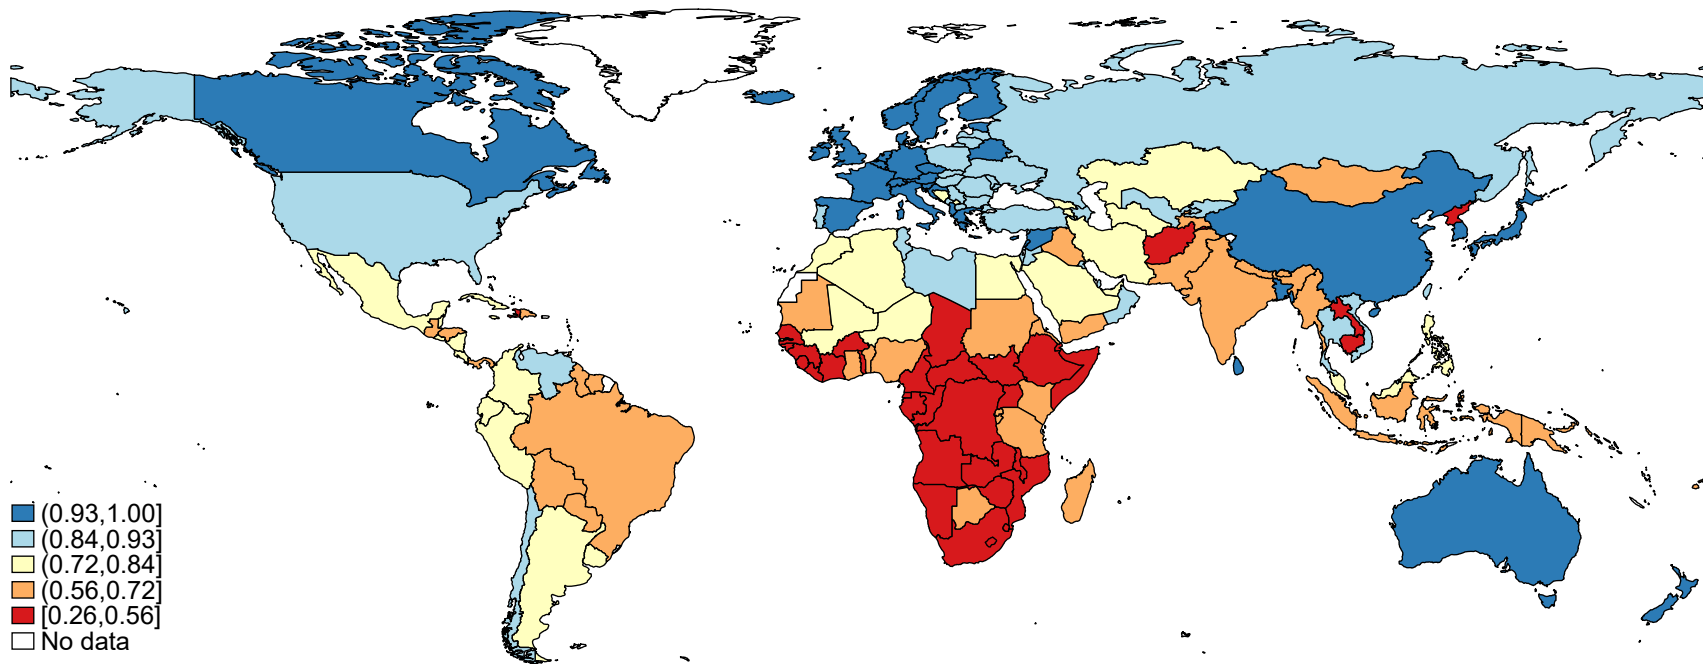
Notes. The map represents the share of total government expenditure on public goods received by the bottom 50% in each country.

Figure A.5.43 – Aggregate Public Education Productivity Around the World, 2019



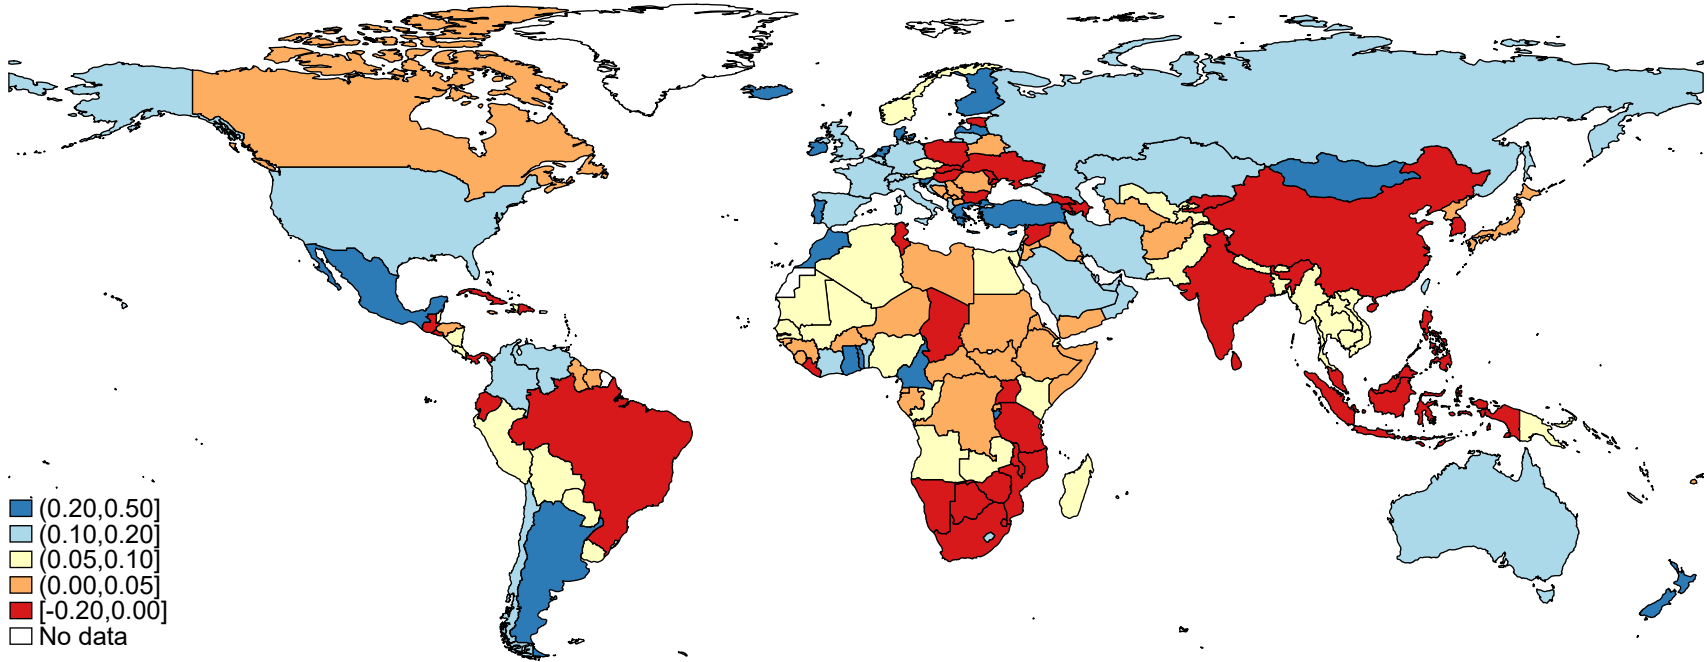
Notes. The map represents estimates of aggregate public education productivity Θ^j in 2019, estimated using public education spending as the only input.

Figure A.5.44 – Aggregate Public Healthcare Productivity Around the World, 2019




Notes. The map represents estimates of aggregate public healthcare productivity Θ^j in 2019, estimated using public health spending as the only input.

Figure A.5.45 – Change in Aggregate Public Education Productivity Around the World, 1980-2019



Notes. The map represents the percentage point change in aggregate public education productivity Θ^j between 1980 and 2019 around the world, estimated using a single-input estimate for each function of government.


Figure A.5.46 – Change in Aggregate Public Healthcare Productivity Around the World, 1980-2019



127

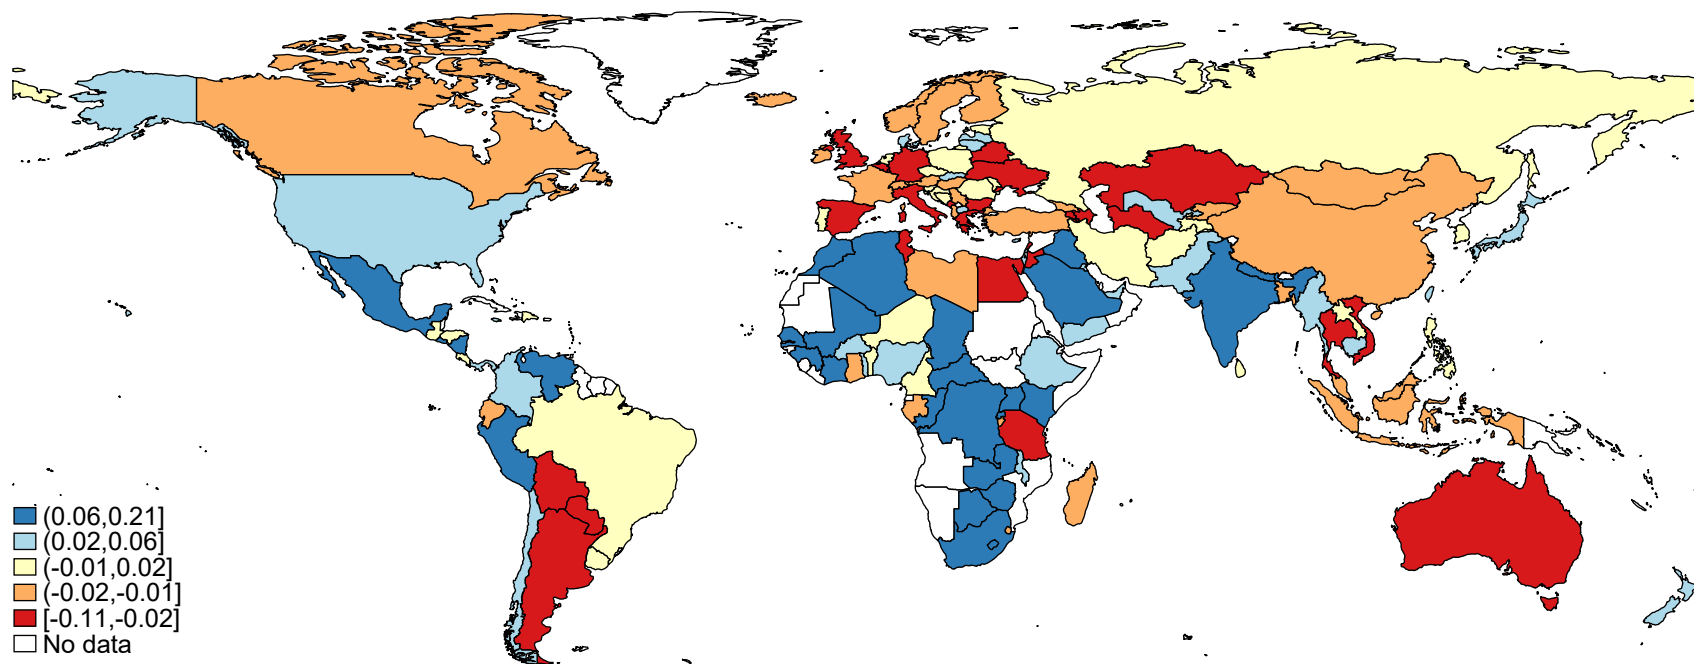
Notes. The map represents the percentage point change in aggregate public healthcare productivity Θ^j between 1980 and 2019 around the world, estimated using a single-input estimate for each function of government.

Figure A.5.47 – Inequality in Public Service Delivery Around the World



Notes. Author's computations using Gallup World Poll data. The figure represents the relative quality of public services received by the bottom 20% of income earners in comparison to the overall population. Values lower than 1 mean that the bottom quintile receive services of lower quality; values higher than 1 mean that they receive services of better quality. Quality is measured as the share of respondents who declare being satisfied with public services in the city or area where they live, for the following services: public transportation systems, roads and highways, the educational system or the schools, the quality of water, and the availability of quality health care. The quality of police services is measured as the share of respondents who declare having confidence in the local police force. These indicators are then aggregated by income quintile, and the ratio of the bottom quintile to the overall average is computed. Finally, the average of this indicator over all public services is calculated, over the entire 2009-2019 period, and represented in the figure.

Figure A.5.48 – Trends in Equal Access to Public Services Around the World, 2009-2019



Notes. Author's computations using Gallup World Poll data. The figure represents the change in the relative quality of public services received by the bottom 20% of income earners, in comparison to the overall population, between 2009-2013 and 2016-2019. Values higher than zero mean that public services have become more progressive; values lower than zero mean that they have become more regressive. Quality is measured as the share of respondents who declare being satisfied with public services in the city or area where they live, for the following services: public transportation systems, roads and highways, the educational system or the schools, the quality of water, and the availability of quality health care. The quality of police services is measured as the share of respondents who declare having confidence in the local police force. These indicators are then aggregated by income quintile, and the ratio of the bottom quintile to the overall average is computed. Finally, the average of this indicator over all public services is calculated over the 2009-2013 and 2016-2019 periods, and the difference between the two periods is represented in the figure.

