Measuring Inequality

Thomas McGregor, Brock Smith and Samuel Wills*

January 2019

Abstract

Inequality is important, both for its own sake and for its political, social and economic implications. However, measuring inequality is not straightforward, as it requires decisions to be made on the variable, population and distributional characteristics of interest. These decisions will naturally influence the conclusions that are drawn so they must be closely linked to an underlying purpose, which is ultimately defined by a social welfare function. This paper outlines important considerations when making each of these decisions, before surveying recent advances in measuring inequality and suggesting avenues for future work.

Keywords: Inequality, poverty, income measurement, social welfare, gini coefficient, nighttime lights, machine learning.

JEL codes: D63, D31, I3, I14, I24, P36, O15

*McGregor: International Monetary Fund; Oxford Centre for the Analysis of Resource Rich Economies, University of Oxford. Smith: Department of Agricultural Economics and Economics, Montana State University. Wills: School of Economics, University of Sydney; Oxford Centre for the Analysis of Resource Rich Economies, University of Oxford; Centre for Applied Macroeconomic Analysis, Australian National University. We thank Kyle Musser for excellent research assistance.

1. *Introduction*

Inequality has been the subject of a lot of popular and academic attention in recent years in the wake of the financial crisis, asset prices elevated by extraordinary monetary stimulus, and growing populist movements around the world (Atkinson and Piketty 2010; Piketty, 2013; 2015; 2018; Milanovic, 2006; Ravaillon, 2018). While inequality is important, it is also a broad concept which makes measuring it a challenge. This paper addresses the key considerations that must be made when measuring inequality, before highlighting recent innovations in data sources and methods used in doing so.

Inequality is a property of a variable's frequency distribution within a population, which is typically summarised in a single statistic. Measuring inequality is therefore not straightforward because it first requires answering a series of questions: what variable do we care about? What population is the focus? And what properties of that variable's distribution matter for our purposes, which we can summarise in a statistic? The answers to each of these questions will depend on the researcher's purpose: there should be different measures, of different variables, for different goals.

The first challenge in measuring inequality is choosing the variable of interest. Inequality is typically thought of with respect to wealth or income. Section 2 shows that there are many other possibilities. Happiness and utility relate more directly to human experience. Utility in particular can be defined to incorporate a wide range of influences on human welfare, including health, social and political outcomes, though is difficult to measure. Wealth is a useful proxy for welfare, and financial wealth is relatively easy to measure (despite difficulties in tracking the wealth of the richest: Zucman, 2013). However, financial wealth excludes components like human capital which, if ignored, can lead to perverse conclusions. Income inequality is also relatively easy to measure, but income varies significantly over a lifetime and there are important distinctions between its pre- and post-tax measurement. Consumption is the main determinant of utility in most economic models, but measurement relies on costly surveys and is the outcome of endogenous decisions and so is not a clean measure of one's circumstances. In all cases, there are important distinctions between inequality of opportunity and inequality of outcome.

The second challenge is in choosing the population within which inequality is measured, as discussed in Section 3. Global inequality does this across all individuals around the world, which is often justified on moral grounds (Singer, 1995). Between-country inequality uses countries, rather than individuals, as the units of observation and is often motivated by countries being the actors in global diplomacy. Finally, within-country inequality studies individuals, or other subgroupings like genders or races, within a nation-state. This is useful as salience is an important aspect of inequality, and most redistributive policies occur at a national level.

The third challenge in measuring inequality is on deciding what properties of the frequency distribution are the focus, and in turn what statistics best summarise them. These properties will depend on the "social-welfare function", the way we choose to aggregate the welfare of all the individuals in the population. Section 4 discusses how some distributions can be ranked while only imposing very general conditions on the social welfare function. However, in most cases ranking levels of inequality involves tradeoffs: often between having a larger pie or dividing it more equally. These tradeoffs are the reason why measures of inequality are not purely statistical, they necessarily embody some social judgement about how to weigh welfare at different points of the distribution.

Inequality matters, both for its own sake and for its political, economic and social implications. These implications – the ultimate reasons why we are interested in inequality – will determine how we choose to measure it. Politically, inequality in wealth and income underpins many key decisions in public policy about how taxes are raised and spent, because democracies must appeal to the median voter (Moffitt et al., 1998), and rulers want to avoid instability (Alesina and Perotti, 1996) and regime change (Gallagher and Hanson, 2009). High inequality is also associated with higher rates of political polarization, which can lead to instability (Payne 2017). Caution is advised here. There is growing evidence that the average person misperceives the level of inequality in their society, underestimating the level of both income and wealth inequality (Hauser and Norton, 2017; Kiatpongsan and Norton, 2014).

Inequality also affects the economy in aggregate. Global inequality in spending power limits potential export markets (Suwa-Eisenmann and Verdie, 2007), in health outcomes increases the risk of global pandemics (Garrett, 2007), and in incomes reduces geopolitical stability (Adelman, 2007). These costs are often disparate and borne in the future, which can make them difficult to address. Within-country inequality in opportunities reduces economic growth by reducing the available stock of human capital (Alesina and Perotti, 1996; Deininger and Squire, 1998; Atkinson et al., 2011), often due to credit market failures that limit opportunities for productive investment (Bourguignon, 2006). Unequal geographic distribution of economic activity can reduce productive matches between workers and jobs (Ioannides and Datcher Loury, 2004). Inequality in wealth can also increase crimes against property (Kelly, 2000), and in income can reduce aggregate demand due to the higher marginal propensity to consume at the bottom of the income distribution (Jappelli and Pistaferri, 2014). Left unchecked this inequality can become entrenched, through inheritance, assortative matching in the marriage market (Gould et al., 2008), political capture and rent seeking (Banerjee et al., 2001) and focusing policy and culture on meritocracy (Young, 1958). On the other hand, wealth inequality may lead to higher savings and investment if wealthy people have a higher propensity to save (Kaldor, 1957), and, if starting at low levels of inequality, may incentivize entrepreneurship (Lazear and Rosen, 1981). Benhabib (2003) develops a theoretical model with a U-shaped relationship between inequality and growth, with positive relationship at lower inequality levels and negative at higher levels.

Socially, inequality tends to reduce the welfare of everyone in society after controlling for wealth or income (Payne, 2017). For those at the bottom of the distribution in unequal societies this includes higher obesity, stress, drug and alcohol dependence, shorter lifespans, and riskier behaviour (Payne, 2017). This all stems from the way that people make decisions based on their status and position relative to others.

There has been growing interest in inequality in recent years, and this has partly been helped by rapid improvements in the tools available to measure it[1]. Section 5 outlines more recent innovative sources, including day and night-time satellite imagery, mobile phone data, and other more speculative sources. Section 7 concludes.

2. *Inequality of what?*

2.1 Opportunity

First, a distinction should be made between equality of opportunity and equality of outcomes, as has previously been done in the study of moral philosophy (e.g. Williams, 1962; Arneson, 2002). Equality of opportunity has much in common with “democratic equality”, or “equality of freedoms” (Anderson, 1999). Equality of opportunity is not opposed to inequality in outcomes, but rather refers to the ability for individuals to freely and fairly generate income, wealth, utility or happiness. It is typically perceived

[1] See appendix section A1 for a discussion of existing data sources for measuring inequality.

to be more desirable than equality in any specific outcome itself. In addition, for a certain level of opportunity, some inequality in outcomes, such as wealth, may even be desirable as a necessary element in the allocation of capital and labour (Illsley, 2002). It is also worth noting that there need not necessarily be any relationship between inequality of opportunity and inequality of outcomes.[2]

Measuring inequality of opportunity in a society is difficult and is typically done by comparing ex-post realised outcomes (such as income) controlling for an individual's endowments (such as parents' socioeconomic status, gender, or ethnic background). Any inequality in income or wealth between individuals that is explained by endowments would point towards the existence of inequality of opportunity. Research by the IMF (2017) finds that inequality of opportunity is higher, on average, in emerging markets, especially in Latin American countries, than in advanced economies (Figure 1). Bhattacharya et al. (2017) take a novel approach to studying equality of opportunity in academic admissions, by identifying the academic performance of the marginal admitted student.

Figure 1: Inequality of opportunity

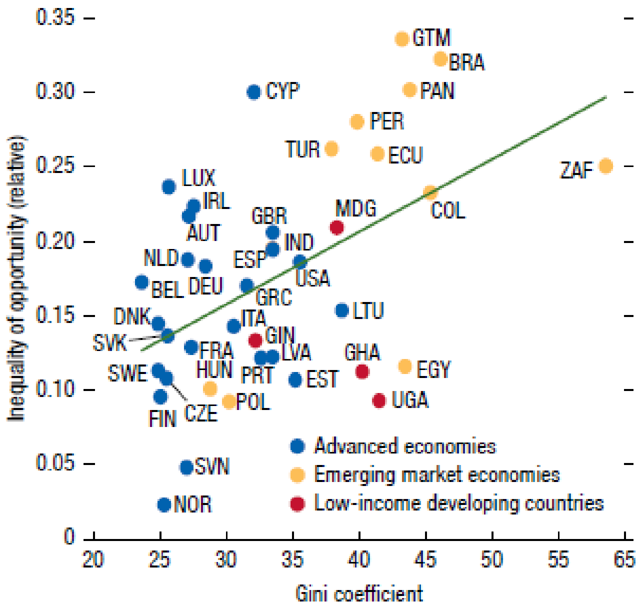

Source: IMF fiscal monitor (2017). Brunori, Ferreira, and Peragine (2013)

Note: Gini coefficients are as of 2015 or the most recent year available. Inequality of opportunity (relative) measures the extent to which circumstances beyond an individual's control (such as family background, gender, and race) affect joint distribution of outcomes (income). It is a lower-bound estimate, because it is not possible to take into account all external circumstances (see Brunori, Ferreira and Peragine, 2013 for details).

However, there are challenges to focusing on equality of opportunity. First, it is not immediately obvious which variables or individual characteristics can be considered "endowments". Many individual traits that lead to higher economic wealth, such as intelligence or determination, may have a large inherited or genetic component. Can these really be ignored as endowments? Then there is the added complication of how to deal with external conditions, ranging from aggregate economic and technological conditions, to parents' wealth and the safety of the neighbourhood where one

[2] To see how consider a society in which the relative position of incomes of each individual in society is defined by random chance. In this society, there is complete equality of opportunity, because all individuals have the same chance of being at the top of the distribution. However, this society is consistent with any actual distribution of outcomes consistent.

grows up. Relative to income and wealth, these factors are inherently difficult to measure in a systematic way. Finally, societies that do strive for equality of opportunity (or "meritocracy") may create self-reinforcing cycles of merit and power that is considered "earned", which can ultimately lead to a more dangerous and polarised society (Young, 1958).

2.2 Happiness and Utility

The economics profession is concerned primarily with welfare, which offers the obvious suggestion of focusing directly on inequality in individuals' well-being or "happiness". A number of studies use self-reported life satisfaction surveys and find that happiness inequality has fallen steadily over time in many countries (Clark, Fleche, and Senik, 2015; Stevenson and Wolfers, 2008).

However, the extensive literature on subjective well-being concludes that i) it is hard to measure—relying on surveys that are subject to bias (Alesina et al., 2004); ii) it is unclear how well-being responds to changes in people's circumstances (Easterlin, 1974; Tella et al., 2003); iii) it is highly correlated with more objective measures such as income (Howell and Howell, 2008; Jebb et al, 2015); and iv) as happiness is non-rivalrous we may be more concerned about the level of happiness, rather than the relative distribution.

Due to the shortcomings in studying happiness, economists often appeal to a similar, but distinct concept: "utility". Utility is a general-purpose measure of welfare. However, there is a circularity in defining it: we care about utility because it is a measure that captures everything we care about. The advantage of utility is that is can be defined very broadly, to be a function of economic variables like income, consumption, wealth and investment opportunities; biological variables like health and life expectancy; social variables like status, friendships and access to the marriage market; environmental variables like access to clean air, water and food; or political variables like ability to vote or access to power. Of course, none of these variables fully capture the richness of people's lives.

It is also difficult to measure utility directly, not least because of the difficulties in defining it. To get around this problem, economists assume that individuals are rational agents who make choices which maximize their own utility given their preferences and available resources. As a result, we can measure utility indirectly by simply measuring the amount of resources available to an individual. A common definition of utility is that in any given period it is monotonically increasing in consumption, and possibly wealth, and monotonically decreasing in hours worked. Therefore, if someone has more resources they will maximise their marginal utility, either by consuming more or working less.

2.3 Consumption

Many studies on poverty and inequality, particularly in developing economies, have focused on consumption or expenditure (e.g. Ravallion, 1995; Chen and Ravallion, 2010). The benefit of such a focus is that consumption enters directly into most definitions of "utility" and so is a relatively direct measure of individual welfare. Consumption is also a relatively stable measure, particularly compared to income, as people theoretically smooth consumption over their lifetime (see permanent income hypothesis (PIH) of Friedman, 1957).

There are of course a number of challenges to using consumption when studying inequality. It is more difficult to measure than income, as it requires tracking expenditures on all goods and services which often means collecting survey data—which are both expensive and time consuming. There is also evidence of misreporting or sample selection in surveys: per capita household consumption from national accounts tends to be substantially higher than average consumption recorded in surveys (Deaton 2005). In addition, consumption inequality may still miss important elements, such as the use or access to public goods (which can be very substantial in some countries).

This then leads us to the most common proxies for utility used in the study of inequality: income and wealth. While neither of these variables capture the full richness of human existence, they are useful in expanding one's opportunity set which makes increasing utility much easier. They also have the advantage of being relatively easy to measure. However, there are many ways in which using income or wealth as a proxy for utility or wellbeing breaks down, including: incomplete markets, capital constraints, information asymmetries, and externalities. In addition, the literature finds mixed evidence on the relationship between income and subjective wellbeing (Benjamin et al., 2012; Kimball et al., 2006).

2.4 Income

Income gives a relatively clear indication of the ability of someone to meet their material needs in the short term. Data on income with long histories is readily available in many countries as a by-product of income taxation (Auten and Gee, 2009; Saez, 2004), and are relatively easy to collect using census or surveys. For these reasons, income has been the main focus of the literature on inequality.

However, income inequality is also far from perfect. One important reason is because income varies substantially over a lifetime: typically increasing until retirement and falling thereafter. Thus, a blind focus on income inequality will have a large intergenerational component: combining the young (low-income and asset-poor) with the retired (low-income and asset-rich). It will also obscure the large degree to which consumption is smoothed over the lifetime – insulating it from income fluctuations. Furthermore, when accumulated wealth can be passed between generations income inequality will give an incomplete picture of the population's actual access to resources.

The type of income matters. Pre-tax inequality reveals important information about the distribution of endowments of skill and wealth, the structure of the labour market, the industrial bargaining and unionisation system, the extent of rent accumulation, and the distribution of financial wealth and thus dividend income. In contrast, post-tax income better reflects the actual resources that people have at their disposal. The choice of measure will depend on the questions at hand.

The source of income is also important. Income derived from labour is inherently different from that derived from the passive ownership of capital, as illustrated in the common utility function that decreases in hours worked as described above. The reason, of course, comes down to the one resource that is perfectly egalitarian: the number of hours in a day. For a given level of income, the more time spent on labour reduces the amount of time available for all other sources of utility.

A distinction can also be made between income that is compensation for labour and capital's marginal product, opportunity cost and risk, and that which can be attributed to the accumulation of economic rents. Income inequality from the former can be an important signal for allocating capital and labour to their most productive uses. However, inequality from the latter can reduce the economy's productive capacity, as labour and capital compete for rents or are employed in unproductive roles that in themselves reduce happiness (see the literature on "bullshit jobs" by Graeber, 2013), rather than pursue more productive monetary or non-monetary activities.

Finally, measures of income inequality will inevitably face reporting issues. First, income usually refers to formal wage income, and income from other sources (like informal wage earnings, rental income, returns from assets, in-kind government assistance, etc.) is often excluded. Second, there is evidence of under-reporting of top incomes in survey data (Alvaredo, 2010). Recent work on the size of top income shares relative to the rest of the distribution has used tax records data (Atkinson and Piketty, 2007; 2010). This avoids missing high income individuals, but still suffers from misreporting, legal or otherwise.

2.5 Wealth

Wealth inequality is also commonly studied as, like income, it measures an individual's access to resources. This has many merits. Financial wealth not only reflects differences in income over a long time span, but also differences in savings rates, inheritances and bequests. While it can be illiquid in the short term, it can be mobilised in the long term. Also, wealth may be more stable than income, and less reliant on personal effort. Thomas Piketty's book "Capital in the Twenty-first Century" (Piketty, 2013), presents a detailed synthesis of evidence on income and wealth inequality in advanced economies, and has sparked a growing debate on these topics.

There are a number of challenges, however, in using wealth to study inequality. First, it is illiquid, so it may not give a clear reflection of one's immediate access to resources. Second, the easily measurable components of wealth – like financial wealth – are incomplete. Take, for example, a student. In many cases their student debt will outweigh their other financial assets, and so they will have negative financial wealth. However, we would not typically say that they are in a worse position than never having studied at all. The reason is because the student has taken on debt to accumulate human capital, which is valuable as a means to increase lifetime earnings but not typically counted as wealth. Policies designed to reduce wealth inequality could therefore have many undesirable consequences, not least redistributing away from people with zero assets towards those with negative assets. These issues could be addressed by valuing human capital, but this is difficult to do. And if we are going to include university education as an asset, then what about other forms of human and social capital, like charisma, coordination, health or one of many other personal assets that can have material worth? Or collective assets, like public schools, hospitals, pension funds and sovereign wealth funds? Ultimately a comprehensive calculation of someone's wealth is very difficult to measure.

Third, the difficultly in measuring wealth inequality is exacerbated by the availability of data. There are many different sources of data on wealth including: national accounts data, micro-level data from household surveys, financial institution surveys, administrative records, tax returns data, and specialised databases on works of arts, luxury cars, etc. There are different challenges associated with the different data sources, including: data quality, comparability across jurisdictions and time, confidentiality rules, response bias and misreporting or mis-sampling.

Finally, life-cycle dimensions lead to different distributions of wealth between generations, with the young typically facing higher levels of indebtedness, and the role of inheritance strengthening the cycle. (Piketty, 2013; Murtin and Mira d'Ercole, 2015).

Wealth inequality appears to be on the rise in many countries (Kopczuk and Saez, 2004; Saez and Zucman, 2016; Piketty, 2014, Piketty and Saez, 2003). Piketty and Zucman (2014) use data on national balance sheets to track how the wealth-to-income ratios have changed over the long run in in 8 advanced economies[3]. They find that wealth-to-income ratios fell in the interwar period, but have since risen steadily in every country back to the levels observed in 18th and 19th century Europe (figure 2). Rognlie (2016) finds evidence that the increased share of capital in total income is mainly due to house price appreciation. Zucman (2019) likewise finds that wealth inequality has risen significantly in the US and at the global level, and argues that these increases may still be underestimated since financial globalization has increased the difficulty of measuring offshore wealth for the top 1%.

Evidence from the OECD Wealth Distribution Database suggests that wealth is much more unequally distributed than incomes, due to the unequal distribution of financial assets (Murtin and Mira d'Ercole, 2015). On average, the top 10% of households holds around 50% of total household net wealth,

[3] US, UK, Germany, France

compared to 25% of total household income. Based on the alternative measure, mean net wealth is 2.5 times larger than median net wealth in the OECD, on average, ranging from 7 times in the United States to around 2 times in the more 'equal' OECD countries.

Figure 2: Private Wealth to National Income Ratios In Europe , 1870-2010

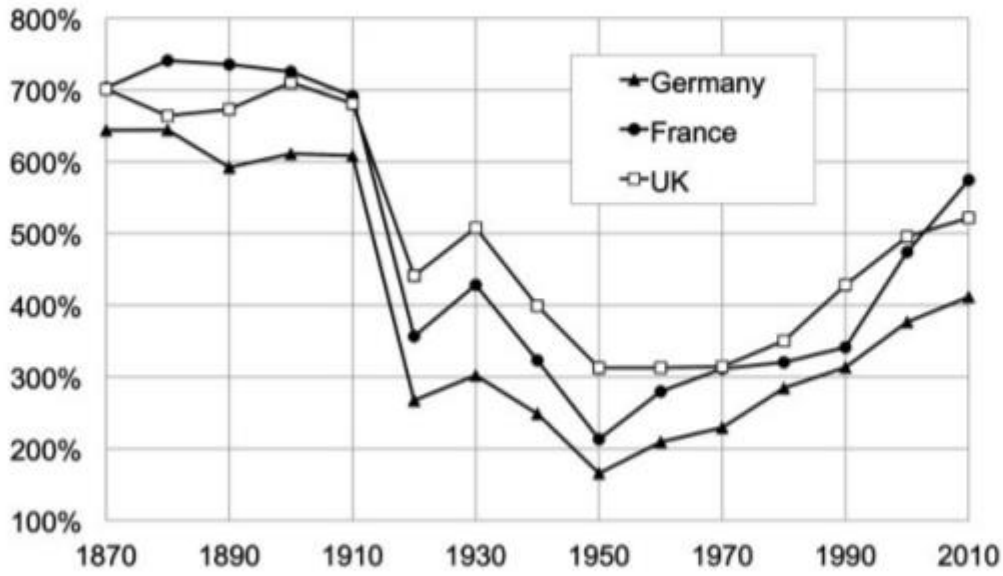

Source: Piketty and Zucman (2014).

Note: Computations using country national accounts. Private wealth = non-financial assets + financial assets - financial liabilities (household and nonprofit sectors). Data are decennial averages (1910–1913 averages for 1910 Europe).

3. *Inequality Between Whom?*

Measuring inequality involves describing the distribution of resources within a group of individuals. The choice of the group, as well as the individual unit, is an important one. Typically this means studying the distribution of resources: across all individuals in the world, between countries, or within countries. The challenge in computing these different inequality measures is finding available and reliable data. Studies on inequality often use combinations of data from different sources (e.g. national accounts vs household survey data).

Any comparison of individuals across countries will require exchange rates to convert income, consumption or wealth into a comparable unit of measurement. If the law of one-price were to hold, we could simply use market exchange rates. However, this is not the case as countries have different economic structures. Purchasing power parity (PPP) exchange rates can overcome this, but need to take into account issues like differences in tastes, the quality of goods, and the provision of non-market goods and services like home childcare or public services (Deaton and Heston, 2010). The World Bank's (WB) International Comparison Program (ICP) is the largest and most comprehensive attempt to collect global price data and compute PPP exchange rates, covering 138 countries over up to 39 years. This addresses quality differences by using very precise descriptions of goods, to ensure like is compared to like, allowing the comparison of living standards across both time and country.[4]

3.1 Global Inequality

Global inequality refers to inequality among the world's 7.6 billion people, regardless of nationality. Milanovic (2012) and Lakner & Milanovic (2013; 2016) find that the global Gini index fell only slightly from 76.3 in 1988 to 75.9 in 2008. In addition, they measure how incomes have changed over time in different parts of the distribution, which gives rise to the well-known "elephant chart" (figure 3).

Figure 3: The Elephant Graph: Global real income change, 1988-2008 (in percent)

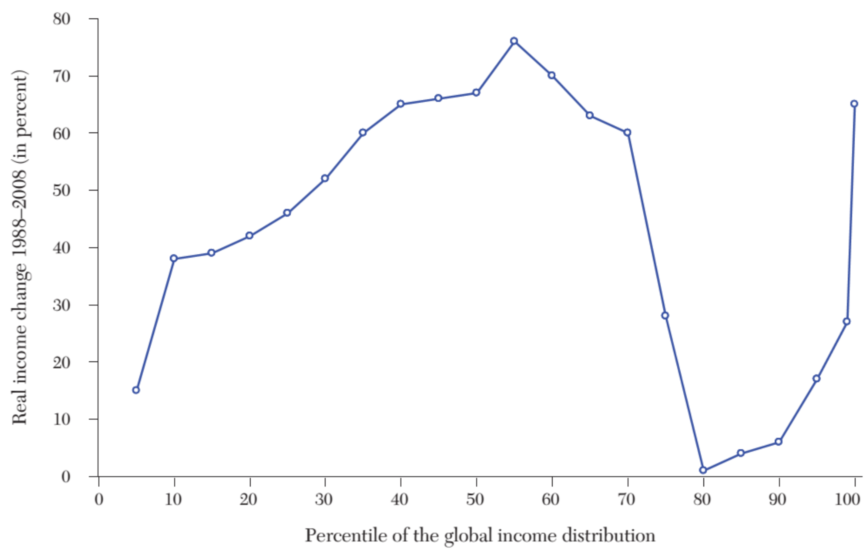

Source: Lakner and Milanovic (2016)

The elephant chart shows relatively slow growth in incomes of the poorest people in the world from 1998-2008, but a sharp increase in incomes for the "rising middle classes" in emerging markets like India and China. Incomes of those at the lower end of the income distribution in the developed world (around the 80th percentile) have grown much more slowly over the same period, while the incomes of the richest have grown by close to 70 percent—which accounts for the vast majority of absolute gains (Ravallion, 2018). Nevertheless, the high growth near the bottom and middle of the global income distribution means that, by many measures, global inequality has fallen (Bourguignon & Morrisson, 2002).

A global definition of inequality has many supporters. Singer (1995) advocates focusing on this type of inequality on moral and philosophical grounds: in essence, one's moral obligation to other humans should not be diminished by their race, nationality or location. There are also many policy issues that naturally lend themselves to focusing on a global definition of inequality. War, pandemics, pollution, global warming, and natural and cultural degradation are all issues of global public goods, and addressing them requires considering global differences in wealth, income, health and regulation, and often necessitate global cooperation and transfers.

There are, however, a number of shortcomings to using global inequality measures. First, it is not clear that world is the natural level at which to consider redistribution, as most of this power is vested in

[4] A complete discussion of comparability of inequality estimates across countries and time is beyond the scope of this paper, but Solt (2009) provides a good overview.

national governments. Furthermore, many of the social and psychological costs of inequality, as opposed to poverty, are positional (Frank, 2008). As a result they rely on salience, which diminishes with distance and exposure. Thus, someone in the middle of the global distribution who has seen their incomes rise rapidly may be less concerned about the positional aspects of their inequality if they are the richest person in the village. Naturally global integration through films, music, and social media are diminishing this, but the point remains. Finally, comparing incomes across countries raises a variety of difficulties, as discussed at the start of this section.

3.2 Between-Country Inequality

An alternative is to consider the world's 200+ countries as the population of interest, and focus on between-country inequality. Countries can either be counted based on their mean income or wealth; or be weighted by their population, though the former is typically the focus. Global inequality rose between 1820 and 1990, driven by rising between-country inequality as today's rich countries developed. This was partially offset by stagnant or falling within-country inequality (Milanovic, 2016; Bourguignon, 2016). Over the past 30 years this pattern has reversed, with the falling global inequality over the past 30 years described above being driven by a decline in between-country inequality (Ravallion, 2018), partially offset by rising within-country inequality.

An advantage of studying between-country inequality is that countries are the units which engage in international diplomacy. Thus, the motives of individual countries in addressing global issues often depends on their position in a between-country distribution, whether it be of income, malaria exposure, carbon emissions or elevation above sea-level. Between country inequality is also one of the main arguments in support of international development assistance.

The disadvantage of focusing on between country inequality is that countries are constructs and do not have utility functions *per se*. Ultimately what matters is the utility of individual people, which is masked in broad groupings. Between-country inequality also suffers from many of the same challenges as global inequality, such as comparing prices between countries and the omission of non-market services.

3.3 Within-Country Inequality

The third common approach is to consider within-country inequality. That is, the distribution of incomes across a society that falls within a national economic boundary (where the individual unit may be persons, households, regions, races, genders, age groups, cohorts, or other groupings). Statistics such as the Gini coefficient, top income shares, and wealth-income ratios have all been used to measure within-country inequality. While between-country inequality has generally fallen, within country inequality, particularly of income, has been growing for the past 30 years—certainly in advanced economies (Lakner and Milanovik, 2013; Dervis and Qureshi, 2016).

It is natural to consider within-country inequality for at least three reasons. First, within country inequality is the most salient, because individuals are surrounded by it every day. Most countries around the world collect data on the distribution of income, wealth and other social, educational and health outcomes, and construct measures of inequality that draw on national consensus. This makes within-country inequality measures arguably the most meaningful in terms of answering important economic political and social policy questions, Certainly at the national level.

Second, countries are typically the level at which most public policy is made. Therefore, any decisions to address inequality, be it through education, health, justice, tax or redistribution tend to happen at the national level. Policy also typically explains why some countries have suffered more from growth in within-country inequality than others (Cornia, 2003).

Third, as taxation is typically levelled at a national level, this is also typically the level at which data is most reliable. However, in emerging markets data still presents a major challenge, which necessitates alternative approaches as discussed in Section 6. Even in the most data-rich developed economies, the picture of wealth and income inequality can be obscured by offshore earnings and company and trust tax structures that are designed to minimise the amount of tax paid. This under-reporting of income is an important consideration when studying inequality((Lakner and Milanovic, 2013; Zucman, 2013; 2019).

Measuring inequality at the national level allows each society to define its own measure and to design policies to tackle the negative elements of inequality that matter most for the citizens for that country. For this reason, it is an extremely using way to measure inequality. However, it does make comparing the level of inequality between countries difficult, particularly seeing as there is no universally agreed approach.

4. *How to measure inequality*

As discussed in the introduction, measuring inequality involves constructing a single statistic that summarises the distribution of a variable within a population, which is agnostic to what the ideal distribution of that variable may be (Kuznets, 1953). However, in most cases, inequality measures are not, and cannot be, purely statistical concepts. Measuring inequality often requires some assumptions to be made regarding how undesirable inequality is in certain parts of the distribution. In more technical terms, inequality measures require some assumptions to be made about the social welfare function.

Simple rankings of income distributions can sometimes be made under very general conditions on the social welfare function. As Atkinson (1970) shows, if we assume that the social welfare function is additively separable and symmetric in individual welfare, which is increasing and concave in income or consumption, then a necessary and sufficient condition for ranking two distributions with the same mean is that one can be obtained from the other by redistributing income from the richer to the poorer—or alternatively, that their Lorenz curves do not intersect.

However, in most real-world cases this condition is not satisfied, and the measurement of inequality involves tradeoffs. The desirability of equality in certain parts of the distribution must be balanced against the incentives to engage in productive activity created by some degree of inequality.

The use of different inequality measures will embody different social judgements about the weight attached to incomes at different points of the distribution. These judgements can either be implicit (e.g., Gini Coefficient) or explicit (e.g., Atkinson Index).

4.1 Graphical approach

One way to describe the level of inequality within a society is through graphs. Various methods have been used in the literature, including histograms, PDFs, CDFs, quintile graphs (Pen, 1971), and Lorenz curves. Graphical presentations can often be more powerful that formulas and statistics in communicating concepts like inequality. It also allows, or forces, the reader to make their own welfare judgments. Figure 4 below, gives examples of some graphical presentations of income distributions.

Figure 4: Graphical presentations of income distributions

(a) Pen's Parade (Quintile Function) for Expenditure in Vietnam (1993 and 1998)

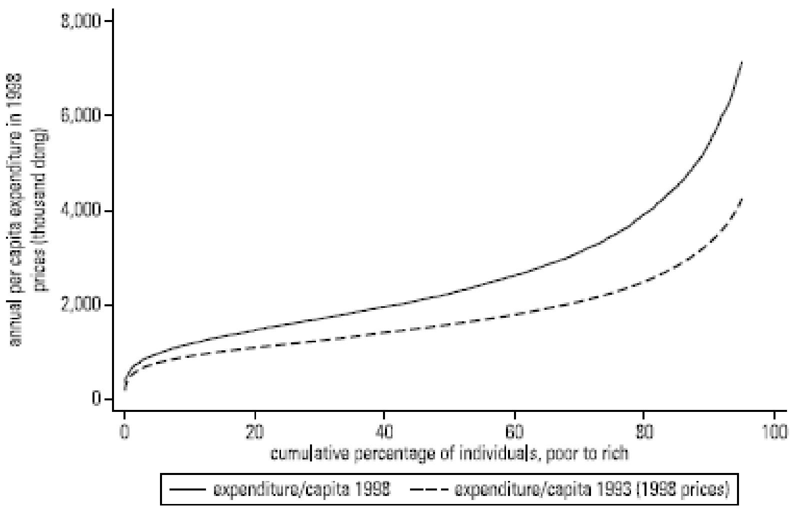

Source: based on data from the Vietnam Living Standards Surveys of 1992–93 and 1998.

Note: This function is truncated at the 95th percentile.

(b) CDFs, Southwest China Income Per Capita (1995 and 1996)

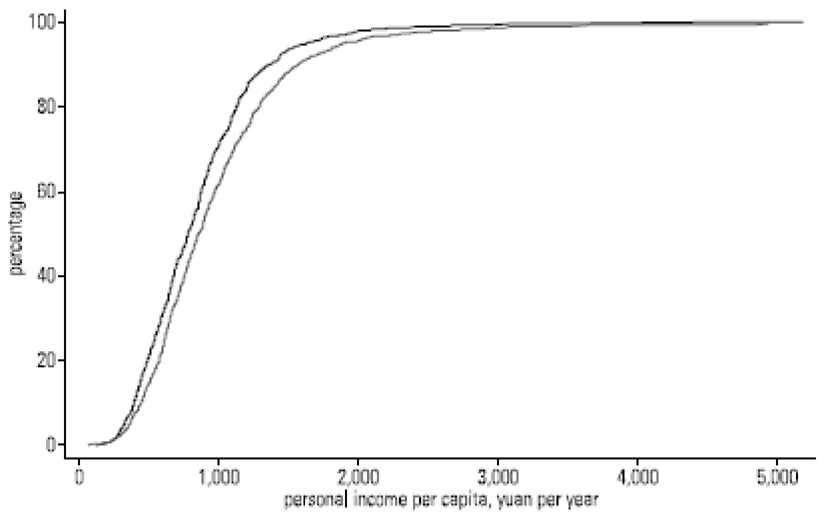

Source: Shaohua Chen, personal communication (reference).

Note: Date for 1995 (upper line) and 1996 (lower line).

Source: Handbook of Poverty and Inequality (Haughton and Khandker, 2009)

There are however two concerns when using a graphical approach. First, we are often interested in a summary description of the level of inequality, which facilitates comparisons across populations. Second, the level of inequality in a society may not be immediately obvious from a graphical presentation. An obvious example is when the Lorenz curves of two societies cross (Figure 5). In these instances, some explicit assumption about social welfare and our preferences for inequality in different parts of the distribution are unavoidable.

Figure 5: When Lorenz Curves Cross

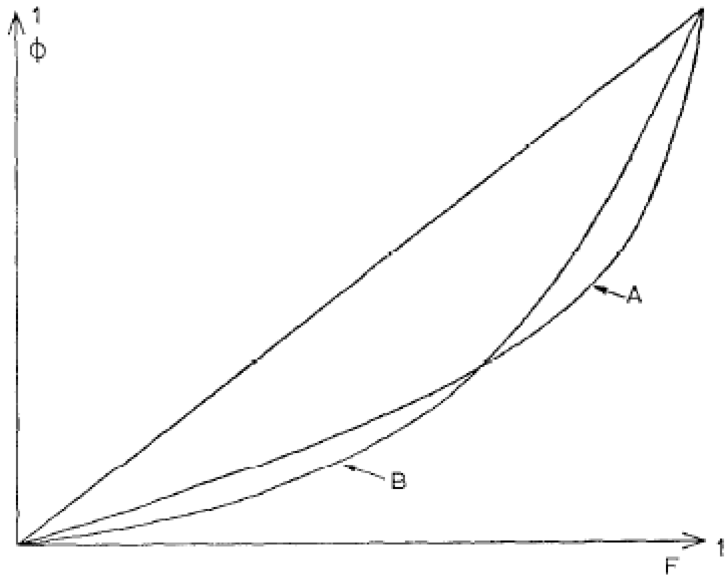

Source: Atkinson (1970)

4.2 Axiomatic approach

The axiomatic approach to inequality measurement brings a more formal structure to designing and comparing inequality measures. It involves setting out and defending *a priori* principles for a desirable inequality measure. The choice of axioms will, in some way or another, be determined by what we think an inequality measure should look like in the first place. Below are six axioms that are generally agreed to be desirable for any measure of inequality:[5]

1. Symmetry/anonymity: any switch of incomes between individuals will leave the measure of inequality unchanged (Fields and Fei, 1978). This is similar to the Rawlsian "veil of ignorance", in that the desired distribution of wealth should be decided on as if one does not know their position in that distribution.
2. Scale invariance: the inequality measure should be insensitive to proportional changes in all incomes (e.g. doubling of incomes or change in the unit of measurement) (Fields and Fei, 1978).
3. Population replication: a proportional change to the underlying population, all else equal, should not change inequality (Dalton, 1920).
4. Pigou-Dalton transfer principle: any transfer of income from one individual to a relatively poorer individual, such that their relative positions do not change, should reduce the measure of inequality (Dalton, 1920).
5. Diminishing transfer principle: a transfer between persons with a given income difference reduces inequality more if these incomes are lower than if they are higher (Kolm, 1976).

Using these axioms, one can assess the suitability of different measures of inequality. Perhaps the most well-known and widely used measure of inequality is the Gini coefficient. The Gini compares the Lorenz curve of a population's income distribution with the line of perfect equality, and can be calculated as follows:

[5] It is worth noting that reasonable alternatives to these axioms are available, and, depending on which axioms one uses, different inequality measures are available.

$$Gini = 1 - \sum_{i=1}^{N} (x_i - x_{i-1})(y_i - y_{i-1})$$

where x_i represents the cumulative measure of individual i's well-being (usually income or expenditure) and y_i is the cumulative population up to individual i. The Gini coefficient is simple to calculate and can be explained using the familiar graphical representation. It satisfies axioms (1)-(4), but does not satisfy axiom (5), as to do so the size of the transfer would need to be large enough to elicit a change in *ranking* of individuals. Criticisms of the Gini index include that it requires implicit assumptions about the form of the underlying social welfare function that are undesirable (Atkinson, 1970), and that it puts too much weight on transfers around the middle of the distribution (Cobham and Sumner, 2013).

Another inequality measure related to the Lorenz curve is the Hoover index (Hoover, 1936), or Robin Hood's index. This measures the proportion of income that would need to be transferred from the top half of the income distribution to the bottom half to achieve maximum equality, and offers a more simple and intuitive measure of inequality than a statistical measure such as the standard deviation. The formula for the Hoover index is as follows:

$$H = \frac{1}{2} \sum_{i=1}^{N} \left| \frac{Y_i}{Y} - \frac{A_i}{A} \right|$$

where Y_i is the income in the i^{th} quantile, N is the number of quantiles, and A_i is the size of the i^{th} quantile. The Hoover index fulfils only axioms (1)-(3), but not the transfer principle—there is no virtue in stealing form the poor to give to the poor.

A useful concept related to the Gini is the inequality extraction ratio (Milanovic 2006). This is the ratio between the maximum possible Gini of a society and the actual Gini. The maximum possible Gini is calculated for a given level of mean income, by providing the whole population with the subsistence level of income, while an infinitesimally small elite capture the remaining share of resources. The locus of maximum Ginis is given by:

$$G(\alpha)^* = \frac{\alpha - 1}{\alpha} (1 - \varepsilon)$$

where α is the average level of income relative to the subsistence level, and ε is the share of the elite in the total population. As $\varepsilon \to 0$ the maximum Gini simplifies to $G(\alpha)^* = \frac{\alpha-1}{\alpha}$, so if the average income level is twice (100 times) that of the subsistence level then the maximum possible Gini would be 50 (99). The extraction ratio is then given by $G/G(\alpha)^*$. Milanovic (2009) uses the extraction ratio for long run comparisons of inequality between countries because it provides a better comparability of inequality between rich and poor societies.

There are a number of inequality measures that satisfy some, if not all, of the six general axioms above. One group of commonly used measures that satisfies axioms 1-5 are the family of generalized entropy (GE) inequality measures, described generally as follows:

$$GE(\alpha) = \frac{1}{\alpha(\alpha - 1)} \left[\frac{1}{N} \sum_{i=1}^{N} \left(\frac{y_i}{\bar{y}} \right)^{\alpha} - 1 \right]$$

where y_i denotes the level of income or expenditure of individual i, $\bar{y}$ denotes the mean over the population, and α is a sensitivity parameter which controls the weight given to individuals at different ends of the distribution—a low value of α makes the index more sensitive to changes at the lower end of the distribution. The GE measure varies between 0 (perfect equality) and infinity.

Another inequality measure that satisfies many of the above axioms is the Theil index (Theil, 1967), which belongs to the entropy measures from information theory (Shannon, 1948). The Theil inequality index is measured as follows:

$$T = \frac{1}{n}\sum_{t=1}^{n} r_i \,.\log\frac{1}{r_i} \qquad r_i = \frac{y_i}{\mu_y} \qquad \mu_y = \frac{\sum_{t=1}^{n} y_i}{n}$$

where r_i is the ratio between the individual income (y_i) and average income (μ_y). The Theil index has the attractive property of decomposability which allows a comparison of inequality between and within different subgroups of the population. It can also be transformed easily into other inequality measures such as the Atkinson index. But perhaps the most famous criticism of the Theil index was provided by Sen (1973) who wrote: "the fact remains that it is an arbitrary formula, and the average of the logarithms of the reciprocals of income shares weighted by income shares is not a measure that is exactly overflowing with intuitive sense."

4.3 *Social Welfare Function Approach*

As previously mentioned, many of the conventional summary measures of inequality make implicit assumptions regarding the form of the underlying social welfare function that are undesirable (Dalton, 1920; Atkinson, 1970). The social welfare function (SWF) approach makes these assumptions explicit. The benefits of this approach are that it provides a complete ranking among alternative distributions and makes explicit the social welfare judgment underlying the measure of inequality. In addition, many SWF-based measures on inequality also satisfy the six axioms discussed above.

These approaches begin by specifying a SWF that we may like to employ in ranking income distributions. One such example is Atkinson's inequality index (Atkinson, 1970)[6] which is based on a constant elasticity of substitution (CES) utility function:

$$W = \frac{1}{N}\sum_{i=1}^{N}\frac{{y_i}^{1-\varepsilon}}{1-\varepsilon}$$

where, y_i denotes individual income, and ε is the inequality aversion parameter[7]. The higher this parameter, the more we care about inequality amongst the poorest in society. The Atkinson index (AI) then gives a measure of inequality as follows:

$$AI = 1 - \left(\frac{1}{N}\sum_{i=1}^{N}\left(\frac{y_i}{\mu}\right)^{1-\varepsilon}\right)^{\frac{1}{1-\varepsilon}}$$

which takes the multiplicative form:

$$AI = 1 - \prod_{i=1}^{N}\left(\frac{y_i}{\mu}\right)^{\frac{1}{n}}$$

Atkinson (1970) goes on to propose comparing income distributions using an "equally distributed equivalent" level of income, in the same way that uncertain incomes are compared using a "certainty-

[6] This is part of a group of SWF based inequality measures that satisfy the axioms proposed by Dalton (1920) Aigner & Heins (1967) and Atkinson (1970), referred to as the Atkinson-Kolm-Sen or A-K-S measures.

[7] When $\varepsilon = 1$ the SWF function collapses to a log specification, as is traditional.

equivalent" measure.[8]The Atkinson index satisfies axioms (1)-(5), by virtue of being based on a SWF that explicitly meets the underlying axioms.

The obvious question that arises when using SWF-based measures of inequality is which SWF we should use, which is important because the ranking of distributions can vary widely depending on the choice. It has been shown, for example, that many 'reasonable' and well-defined social welfare functions result in a different ranking of income distributions, with identical means, than would be obtained by the Gini coefficient (Atkinson, 1970; Dasgupta, Sen and Starrett, 1970; Lambert, 1984).

4.4 Focusing on the Extremes

There are many more inequality measures available, each satisfying a different set of axioms. Describing them all is far beyond the scope of this paper, but one set of measures is worth mentioning. These are measures that focus on specific parts of the income distribution.

When asked, people often reject the transfer principle (axiom 4), preferring instead to think of inequality as an aggregation of distances of each individual from some reference point—often themselves. This has given rise to inequality 'complaints' measures (Temkin, 1986; Devooght, 2003; Cowell and Ebert, 2004), which are typically upward-looking. In a similar vein, Saez, Piketty and Atkinson (2011) track the incomes of the richest people in society in 30 OECD countries, from the richest 10 percent all the way to the richest 0.01 percent, essentially ignoring the vast majority of individual incomes.

The Palma index[9], which is in the family of inter-decile ratios, is the ratio of national income shares of the top 10 percent of households relative to the bottom 40 percent. The motivation for this measure is based on the observation by Palma (2006; 2011) that the income share of the middle 50 percent has been remarkable stable across both time and countries. This means that focusing on the ratio between the two points in the distribution (the 10^{th} and 40^{th} percentiles), should capture the central features of comparative inequality. Cobham et al. (2016) present evidence in developing countries to support the Palma and show that, in practice, this index explains much of the variation in the Gini.

4.5 Unequal inequality

Up until now, all of the inequality measures considered have upheld the first axiom of symmetry or anonymity. While this is an attractive property for a measure of inequality, it may go against societies' natural tendency to care about 'who' is affected by inequality. Consider the following thought experiment. Imagine two societies with the same population size and identical distributions of income. Now let society A be such that the positions of each individual in the distribution never changes. The poor are condemned to be poor and the rich blessed to be rich, both for eternity. Society B on the other had is such that there is some arbitrary lottery such that individuals are assigned their position in the distribution by chance, every few years. Whilst we may not all agree, we almost certainly will all have some view as to which society is more "equal" or "fair". Yet none of the measures that exist today can distinguish between these two societies.

This brings us to a second issue: that of judging policy interventions that lead to a reallocation of resources between individuals or groups in society. Whilst we may all be able to agree that reducing

[8] Another SWF approach is that of Dalton (1920) who takes the ration of actual to the highest possible social welfare.

[9] After Palma (2006; 2011) observed that the income share of the middle 50% has been remarkable stable across time and countries.

inequality is a good thing, we are unlikely to agree on who the specific winners and losers from a policy should be. Consider a proposal for two opposing policies that each reduce inequality for two different groups in society. How do we choose between them? What if one group happens to be a part of society that has been historically disadvantaged, such as ethnic minorities, or women?

5. *New data and methods for estimating inequality*

In Appendix A1 we provide a brief overview of traditional data sources used to measure inequality, including censuses, other large-scale surveys, and tax records. These traditional sources have several flaws. Poorer households are often missed from surveys (Carr-Hill, 2013), high-income households don't respond (Korinek et al., 2005; Anand and Segal, 2015), and income is often undeclared (Zucman, 2013). In rich countries, response rates to income surveys have been steadily dropping since at least 2000 (The Economist, 2018). The sources are also inconsistent over time, which makes statistical analysis difficult. Censuses are at a minimum once every ten years. Country-wide surveys are irregular, and have widely different methodologies, response rates, attrition and coverage. According to Jean et al (2016), from 2000-2010 39 of 59 African countries conducted less than two surveys adequate for estimating poverty rates, with 14 countries conducting none. Tax records are generally not publicly available and limited to wealthy countries. Finally, there may be concepts other than income or wealth that are traditionally difficult to measure but whose distribution within a society is important from an economic standpoint.

In this section we survey some newer measures of inequality that partially address these problems, made possible by recently available data sources and computing techniques. We also discuss the limitations of these measures, and suggest some potential sources exploiting "big data". Many of the sources discussed in this section have particular potential in developing countries, where data on the distribution of income and wealth is weakest.

5.1 Spatial lights and population data

Satellites have been recording data on night-time light intensity around the globe since the early 1990s. This data is highly correlated with economic output within a given area (Elvidge, 1997; Henderson et al., 2011) due to the ambient luminosity from cars, streetlights, and electric lights in households and firms. As a result it has attracted significant attention from economists in recent years.[10]

Lights data has proven useful as a spatially disaggregated measure of income (Weidmann and Schutte 2017, Mellander et al 2015, and Bruederle and Hodler 2018). Similarly, the share of people living in darkness is an effective measure both of rural poverty, and regional inequality (Smith and Wills 2018). However, using nighttime lights to measure individual inequality is more limited. This is because it is difficult to use spatial data for studying densely populated areas. While lights and spatial population data (such as LandScan and Gridded Population of the World) have high resolution relative to traditional sources, in a densely populated city one pixel can still cover tens of thousands of people, including extreme rich and poor alike. Elvidge et al (2012) combine lights and spatial population data

[10] There are two main publicly available global nighttime lights data sets. The US Air Force Defense Meteorological Satellite Program – Operational Linescan Systems (DMSP-OLS) covers the years 1992-2013, and for later years the Visible Infrared Imaging Radiometer Suite (VIIRS) provides more accurate light intensity readings at a higher resolution.

to construct a lights-based Gini coefficient[11], but finds only a weak correlation with traditional income Ginis (R^2 of 0.1)[12]. However, the authors do find strong inverse correlations between the NLDI and various broad development indicators like electrification rates, poverty rates and the Human Development Index, and thus interpret it as more of a general development measure than inequality measure.

However, lights data can yield insights about *regional* inequality. Lessmann and Seidel (2017) use lights and spatial population data to create a global data set of regional Gini coefficients where administrative regions are the unit of analysis rather than individuals. Different regional classifications may also be used. The Global Rural-Urban Mapping Project classifies cells as urban or rural at the same resolution as the DMSP-OLS lights data, allowing for rural vs. urban comparisons. Alesina et al (2016) combines lights and spatial population data with a map of ethno-linguistic homelands in Africa to construct ethnically-based regional Gini coefficients. Minard (2018) similarly constructs regional Ginis based on both administrative regions and ethnic homelands in China. Smith and Wills (2018) use nighttime lights to show that oil discoveries in developing countries increase regional inequality: the additional wealth goes to illuminating cities rather than the rural poor living in darkness.

5.2 Machine learning from satellite images

Machine learning applications have become increasingly common in economics in recent years (Athey, 2017) and present exciting opportunities for processing the vast amounts of geo-spatial data becoming available. A major effort in this vein, and to our knowledge the only attempt as of this writing at using machine learning techniques to measure household wealth, is Jean et al (2016). This paper uses a convolutional neural network (CNN) that is “trained” to predict wealth and consumption in five African countries[13] from high-resolution daytime image data publicly available from Google Static Maps. The CNN extracts identifiable features (e.g. roads, schools, house types, etc) from the image data to predict nighttime light intensity. Then, after having learned which subset of features is useful for predicting lights, the CNN is then trained with these features only to predict survey-based measures of wealth and consumption.[14] Wealth and consumption are estimated at a cluster level, where a cluster roughly corresponds to a rural village or an urban ward.

The resulting model predictions are found to explain 37-55% of the variation in household consumption and 55-75% of the variation in asset wealth across the five countries. The predictions outperform predictions based solely on nighttime lights in 81% of trials, and 99.5% of trials for clusters that are three times below the poverty line. Further, predictions generated from a model trained in one country still perform well in the other sample countries, so out-of-sample predictions would likely be accurate at least in other developing African countries. This paper is largely a proof-of-concept and to our knowledge the technique has not been used to generate a worldwide data set of predicted consumption and wealth levels, but in theory all countries with sufficient survey data could be used to train a model that is then applied to the rest of the world. However, while this approach is suitable

[11] This effectively assigns each person within a pixel that pixel’s average lights per capita value.

[12] Constructing this measure using VIIRS data may yield much better inequality estimates, due to the higher resolution and lack of top-coding.

[13] Nigeria, Tanzania, Uganda, Malawi, and Rwanda.

[14] The CNN does not predict survey-based wealth from images directly because there are not enough wealth observations available until the number of relevant features is substantially reduced in the first step.

to find present-day levels of poverty and inequality, its ability to track changes over time is limited to the extent that year-by-year daytime image data is available. Related work has been used to generate maps of childhood malnutrition (Osgood-Zimmerman et al, 2018)

5.3 Mobile phone metadata

Another novel approach to estimating income and wealth distributions is using information from anonymized mobile phone metadata. Mobile phone data captures information on the frequency of communications, patterns of travel, and users' social structures. Soto et al (2011) and Smith-Clarke et al (2014) are among the earliest to demonstrate the viability of this approach. Blumenstock et al. (2015) compare data on billions of phone interactions in Rwanda's largest mobile phone network to wealth indicators for 856 people collected by phone surveys. They find that 68% of the variation in subscribers' wealth can be predicted by phone data. The most predictive measures are those that indicate the "entropy" of the user's contacts' outgoing calls and messages. In other words, when a user's contacts tend to send communications at consistent times of the day and week, it indicates that the user and their contacts are likely to be professionals with regular work schedules. The trained model is then applied to produce predictions for all mobile phone users in the sample, and these predictions are compared at the cluster level to DHS surveys. The predictions explain 91.7% of the average cluster-level variation of DHS surveys.

A related study by Steele et al (2016) predicts wealth levels at a high resolution in Bangladesh using mobile phone data and several publicly available geo-spatial data sets including night-time lights, soil and climate data, and distance to roads and major cities. Again, compared to geo-coded survey data they find highly accurate predictions ($r^2 = .76$ for the whole country). Further, they find that using only mobile phone data yields comparable accuracy, allowing for higher resolution and frequency of wealth estimates.

Similar to daytime satellite images, the viability of using mobile phone data to conduct statistical studies of inequality will be limited to the extent that it is available across countries and years. But the methodology is easily scalable in theory and is a promising prospect for future inequality data. It also presents an opportunity to reduce inequality, with mobile transactions data being used to evaluate creditworthiness, in turn expanding financial inclusion (Björkegren and Grissen, 2015; Francis et al., 2017)

5.4 Potential "big data" measures

"Google Trends" was released in 2009. This allows users to see how frequently any search term has been used over time. Search data has already proven useful in a variety of applications, including early detection of disease outbreaks (Alessa and Faezipour, 2018) and forecasting tourist inflows (Park et al, 2017). Stephens-Davidowitz (2017) makes the case that due to the perceived privacy of google searches, they reveal true attitudes and preferences of users that they frequently lie about in surveys. Importantly, Google Trends allows users to find search frequencies by regions and cities throughout the world (in the US, the regions are MSAs. Other countries use states or other administrative regions), so the tool can be used to estimate regional distributions, though not individual ones. Stephens-Davidowitz (2014) provided an initial proof-of-concept for the viability of using Google Trends data in economics research, finding that racist search terms within a region were predictive of vote shares for Barack Obama relative to John Kerry. Several studies since have used Google Trends, including Baker and Fradkin (2017), which develops a "Google Job Search Index" (GJSI) to estimate job-searching activity. The GJSI could arguably be interpreted as an example of perceived "opportunity", as it may reflect optimism towards the labour market (at least when holding constant the number of unemployed). Another possible index along these lines is one that measures search activity related to college applications. Further, to our knowledge Google Trends data has not been used to estimate regional income or wealth, but this is a promising prospect for future work.

Social media offers further possibilities for inequality measures. Data on Facebook activity could be thought of as a significantly richer version of mobile phone network metadata that has already proven to be an accurate estimator of wealth. It also offers truly global scope, with around 2.2 billion users as of this writing. Kosinski et al (2013) demonstrated that Facebook likes alone could yield highly accurate predictions of demographic characteristics, as well as traits like political views, intelligence, use of addictive substances, and personality traits. There are legal and ethical issues with these type of data that are beyond the scope of this paper, but if researchers were able to acquire social network activity data that credibly protected user privacy, it would offer exceptional possibilities for producing individual inequality measurements of all types across the world.

6. Conclusion

The issue of inequality has exploded in the popular consciousness in recent years. Yet the causes and effects of inequality remain difficult to study, as there are both conceptual and practical challenges in measuring it. This paper first reviewed the issues that must be addressed when measuring inequality, including deciding on the variable, population and properties of the distribution that are of interest. This shows that there is no “best” single measure of inequality: different measures may emphasize inequality in different parts of the distribution, and thus yield widely different conclusions. Therefore the choices should be tailored to the goals. We then provide an overview of data sources that can be used to measure inequality, with a focus on recent innovative sources that correct for some of the shortcomings of traditional sources, while having some shortcomings of their own. Continuing to improve upon these measures is important as economics shifts its focus from aggregate measures to a more detailed appreciation of distributions.

References

Adelman, C. (2007). Foreign Aid: Effectively Advancing Security Interests. *Harvard International Review; Cambridge*, *29*(3), 62–67.

Aigner, D. J., & Heins, A. J. (1967). On the determinants of income equality. *The American Economic Review*, *57*(1), 175-184.

Alesina, A., Di Tella, R., & MacCulloch, R. (2004). Inequality and happiness: are Europeans and Americans different? *Journal of Public Economics*, *88*(9–10), 2009–2042.

Alesina, A., Michalopoulos, S., & Papaioannou, E. (2016). Ethnic inequality. *Journal of Political Economy*, *124*(2), 428-488.

Alesina, A., & Perotti, R. (1996). Income distribution, political instability, and investment. *European Economic Review*, *40*(6), 1203–1228.

Alessa, A., & Faezipour, M. (2018). A review of influenza detection and prediction through social networking sites. *Theoretical Biology and Medical Modelling*, *15*(1), 2.

Alvaredo, F. (2010). The rich in Argentina over the Twentieth Century, 1932-2004. *Top incomes: A global perspective*, *2*.

Alvaredo, F., Chancel, L., Piketty, T., Saez, E., & Zucman, G. (2018). The Elephant Curve of Global Inequality and Growth. *AEA Papers and Proceedings*, *108*, 103–108.

Anand, S., & Segal, P. (2015). The global distribution of income. In *Handbook of income distribution* (Vol. 2, pp. 937-979). Elsevier.

Anderson, E. S. (1999). What is the Point of Equality?. *Ethics*, 109(2), 287-337.

Arneson, R. (2002). "Equality Of Opportunity". Stanford Encyclopedia Of Philosophy.

Athey, S. (2017). The impact of machine learning on economics. In *Economics of Artificial Intelligence*. University of Chicago Press.

Atkinson, A. B. (1970). On the measurement of inequality. *Journal of economic theory*, *2*(3), 244-263.

Atkinson, A. B., & Piketty, T. (Eds.). (2007). *Top incomes over the twentieth century: a contrast between continental european and english-speaking countries*. OUP Oxford.

Atkinson, A. B., & Piketty, T. (Eds.). (2010). *Top incomes: A global perspective*. Oxford University Press.

Atkinson, A. B., Piketty, T., & Saez, E. (2011). Top incomes in the long run of history. *Journal of economic literature*, *49*(1), 3-71.

Auten, G., & Gee, G. (2009). Income Mobility in the United States: New Evidence from Income Tax Data. *National Tax Journal*, *62*(2), 301–328.

Bagde, S., Epple, D., & Taylor, L. (2016). Does Affirmative Action Work? Caste, Gender, College Quality, and Academic Success in India. *American Economic Review*, *106*(6), 1495–1521.

Baker, S. R., & Fradkin, A. (2017). The Impact of Unemployment Insurance on Job Search: Evidence from Google Search Data. *Review of Economics and Statistics*, *99*(5), 756-768.

Banerjee, A., Mookherjee, D., Munshi, K., & Ray, D. (2001). Inequality, Control Rights, and Rent Seeking: Sugar Cooperatives in Maharashtra. *Journal of Political Economy*, *109*(1), 138–190.

Benhabib, J. (2003). The tradeoff between inequality and growth. *Annals of economics and finance*, *4*, 491-507.

Benjamin, D. J., Kimball, M. S., Heffetz, O., & Rees-Jones, A. (2012). What Do You Think Would Make You Happier? What Do You Think You Would Choose? *The American Economic Review*, *102*(5), 2083–2110.

Bennett, M. M., & Smith, L. C. (2017). Advances in using multitemporal night-time lights satellite imagery to detect, estimate, and monitor socioeconomic dynamics. *Remote Sensing of Environment*, *192*, 176-197.

Berrittella, M. (2017). Can Stability of Foreign Aid Agreement Reduce Global Income Inequality? *Economic Analysis and Policy*, *54*, 105–111.

Bhattacharya, D., Kanaya, S., & Stevens, M. (2017). Are university admissions academically fair?. *Review of Economics and Statistics*, *99*(3), 449-464.

Björkegren, D., & Grissen, D. (2017). Behavior revealed in mobile phone usage predicts loan repayment. arXiv preprint arXiv:1712.05840.

Blumenstock, J., Cadamuro, G., & On, R. (2015). Predicting poverty and wealth from mobile phone metadata. *Science*, *350*(6264), 1073-1076.

Bornschier, V., Chase-Dunn, C., & Rubinson, R. (1978). Cross-National Evidence of the Effects of Foreign Investment and Aid on Economic Growth and Inequality: A Survey of Findings and a Reanalysis. *American Journal of Sociology*, *84*(3), 651–683.

Bourguignon, F. (2006). From income to endowments: the difficult task of expanding the income poverty paradigm. *Poverty and inequality*, 76-102.

Bourguignon, F., & Morrisson, C. (2002). Inequality Among World Citizens: 1820–1992. *American Economic Review*, *92*(4), 727–744.

Bruederle, A., & Hodler, R. (2018). Nighttime lights as a proxy for human development at the local level. *PloS one*, *13*(9).

Brunori, Paolo, Francisco H. G. Ferreira, and Peragine, V. (2013). Inequality of Opportunity, Income Inequality, and Economic Mobility: Some International Comparisons. World Bank Policy Research Working Paper 6304.

Carr-Hill, R. (2013). Missing millions and measuring development progress. *World Development*, *46*, 30-44.

Checchi, D., & Peragine, V. (2010). Inequality of Opportunity in Italy. *Journal of Economic Inequality*, *8*(4), 429–450.

Chen, S., & Ravallion, M. (2010). The developing world is poorer than we thought, but no less successful in the fight against poverty. *The Quarterly Journal of Economics*, *125*(4), 1577-1625.

Clark, A. E., Flèche, S., & Senik, C. (2016). Economic growth evens out happiness: Evidence from six surveys. *Review of Income and Wealth*, *62*(3), 405-419.

Cobham, A., & Sumner, A. (2013). Is it all about the tails? The Palma measure of income inequality.

Cobham, A., Schlögl, L., & Sumner, A. (2016). Inequality and the tails: the Palma Proposition and Ratio. *Global Policy*, *7*(1), 25-36.

Cornia, G. A. (2003). The Impact of Liberalisation and Globalisation on Within-country Income Inequality. *CESifo Economic Studies; Oxford*, *49*(4), 581–616.

Cowell, F. (1977). *Measuring inequality: techniques for the social sciences*. Wiley.

Cowell, F., & Ebert, U. (2004). Complaints and inequality. *Social Choice and Welfare*, *23*(1), 71-89.

Dalton, H. (1920). The measurement of the inequality of incomes. *The Economic Journal*, *30*(119), 348-

361.

Dasgupta, P., Sen, A. & Starrett, D (1973). Notes on the measurement of inequality. Journal of Economic Theory, 6(2), pp.180-187.

Deaton, A. (2005). Measuring poverty in a growing world (or measuring growth in a poor world). *Review of Economics and statistics*, *87*(1), 1-19.

Deaton, A., & Heston, A. (2010). Understanding PPPs and PPP-based national accounts. *American Economic Journal: Macroeconomics*, *2*(4), 1-35.

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

Derviş, K., & Qureshi, Z. (2016). Income distribution within countries: rising inequality, mimeo, Brookings Institution.

Devooght, K. (2003). Measuring inequality by counting 'complaints': theory and empirics. *Economics & Philosophy*, *19*(2), 241-263.

Easterlin, R. A. (1974). Does Economic Growth Improve the Human Lot? Some Empirical Evidence. In P. A. David & M. W. Reder (Eds.), *Nations and Households in Economic Growth* (pp. 89–125). Academic Press.

Elvidge, C. D., Baugh, K. E., Anderson, S. J., Sutton, P. C., & Ghosh, T. (2012). The Night Light Development Index (NLDI): a spatially explicit measure of human development from satellite data. *Social Geography*, *7*(1), 23-35.

Francis, E., Blumenstock, J., & Robinson, J. (2017). Digital credit: A snapshot of the current landscape and open research questions. *Bureau for Research and Economic Analysis of Development Working Paper*, 516.

Frank, R. H. (2008). Should public policy respond to positional externalities? *Journal of Public Economics*, *92*(8), 1777–1786.

Friedman, M. (1957). Introduction to" A Theory of the Consumption Function". In A theory of the consumption function (pp. 1-6). *Princeton university press.*

Fields, G. S., & Fei, J. C. (1978). On inequality comparisons. *Econometrica: Journal of the Econometric Society*, 303-316.

Gallagher, M., & Hanson, J. K. (2009). Coalitions, Carrots, and Sticks: Economic Inequality and Authoritarian States. *PS, Political Science & Politics; Washington*, *42*(4), 667-672,631.

Garrett, L. (2007). The Challenge of Global Health. *Foreign Affairs*, *86*(1), 14–38.

Goedhart, T., Halberstadt, V., Kapteyn, A., & Van Praag, B. (1977). The poverty line: concept and measurement. *Journal of Human Resources*, 503-520.

Gould, E. D., Moav, O., & Simhon, A. (2008). The Mystery of Monogamy. *American Economic Review*, *98*(1), 333–357.

Graeber, D. (2013). On the phenomenon of bullshit jobs: A work rant. *Strike Magazine*, (3).

Haughton, J., & Khandker, S. R. (2009). *Handbook on poverty+ inequality*. World Bank Publications.

Hauser, O. & Norton, M. (2017). (Mis)perceptions of inequality. Current Opinion in Psychology. Vol 18, pages 21-25.

Herzer, D., & Nunnenkamp, P. (2012). The effect of foreign aid on income inequality: Evidence from panel cointegration. *Structural Change and Economic Dynamics*, *23*(3), 245–255.

Hoover, E. M. (1936). The Measurement of Industrial Localization, Review of Economics and Statistics, 18, No. 162–71.

Howell, R. T., & Howell, C. J. (2008). The relation of economic status to subjective well-being in developing countries: A meta-analysis. *Psychological bulletin*, *134*(4), 536.

Illsley, R. (2002). A City's Schools: From equality of input to inequality of outcome. *Oxford Review of Education*, *28*(4), 427–445.

Ioannides, Y. M., & Datcher Loury, L. (2004). Job Information Networks, Neighborhood Effects, and Inequality. *Journal of Economic Literature*, *42*(4), 1056–1093.

Jappelli, T., & Pistaferri, L. (2014). Fiscal Policy and MPC Heterogeneity. *American Economic Journal. Macroeconomics; Pittsburgh*, *6*(4), 107–136.

Jean, N., Burke, M., Xie, M., Davis, W. M., Lobell, D. B., & Ermon, S. (2016). Combining satellite imagery and machine learning to predict poverty. *Science*, *353*(6301), 790-794.

Jebb, A. T., Tay, L., Diener, E., & Oishi, S. (2018). Happiness, income satiation and turning points around the world. *Nature Human Behaviour*, *2*(1), 33.

Kaldor, N. (1957). A model of economic growth. *The economic journal*, *67*(268), 591-624.

Kelly, M. (2000). Inequality and crime. *Review of economics and Statistics*, *82*(4), 530-539.

Kiatpongsan, S. & Norton, M. (2014). How much (more) should CEOs make? A universal desire for more equal pay. Perspectives on Psychological Science. 9:587-593.

Kimball, M., Willis, R., Rangel, A., Rayo, L., Shapiro, M., Silverman, D., Tsutsui, Y. (2006). *Utility and Happiness*.

Kolm, S. C. (1976). Unequal inequalities. I. *Journal of economic Theory*, *12*(3), 416-442.

Korinek, A., Mistiaen, J. A., & Ravallion, M. (2005). *Survey nonresponse and the distribution of income*. The World Bank.

Kopczuk, W., & Saez, E. (2004). Top Wealth Shares in the United States, 1916-2000: Evidence From Estate Tax Returns. *National Tax Journal*, *57*(2), 445-487.

Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 201218772.

Kuznets, S. S. (1953). *Economic change*. Norton.

Lazear, E. P., & Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of political Economy*, *89*(5), 841-864.

Lakner, C., & Milanovic, B. (2013). *Global income distribution: from the fall of the Berlin Wall to the Great Recession*. The World Bank.

Lakner, C., Milanovic, B. (2016). Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession. *World Bank Economic Review*, 30(2), 203-232.

Lessmann, C., & Seidel, A. (2017). Regional inequality, convergence, and its determinants–A view from outer space. *European Economic Review*, *92*, 110-132.

Mejia, D., & St-Pierre, M. (2008). Unequal Opportunities and Human Capital Formation. *Journal of Development Economics*, *86*(2), 395–413.

Mellander, C., et al. (2015). "Night-time light data: A good proxy measure for economic activity?" *PloS one* 10(10).

Milanovic, B. (2006). *Global Income Inequality : What It Is And Why It Matters*. The World Bank.

Milanovic, B. (2009). Global inequality and global inequality extraction ratio: The story of the last two centuries.

Milanovic, B. (2012). Global inequality recalculated and updated: the effect of new PPP estimates on global inequality and 2005 estimates. *The Journal of Economic Inequality*, *10*(1), 1-18.

Minard, P. (2018). Is China's regional inequality ethnic inequality?. Unpublished working paper.

Moffitt, R., Ribar, D., & Wilhelm, M. (1998). The decline of welfare benefits in the U.S.: the role of wage inequality. *Journal of Public Economics*, *68*(3), 421–452.

Murtin, F., & d'Ercole, M. M. (2015). Household wealth inequality across OECD countries: new OECD evidence. Mimeo, OECD.

Osgood-Zimmerman, A., Millear, A.I., Stubbs, R.W., Shields, C., Pickering, B.V., Earl, L., Graetz, N., Kinyoki, D.K., Ray, S.E., Bhatt, S. and Browne, A.J. (2018). Mapping child growth failure in Africa between 2000 and 2015. Nature, 555(7694), 41.

Palma, J. G. (2006). Globalizing Inequality: 'Centrifugal' and 'Centripetal' Forces at Work. DESA Working Paper No. 35.

Palma, J. G. (2011). Homogeneous middles vs. heterogeneous tails, and the end of the 'inverted-U': It's all about the share of the rich. *Development and Change*, *42*(1), 87-153.

Park, S., Lee, J., & Song, W. (2017). Short-term forecasting of Japanese tourist inflow to South Korea using Google trends data. *Journal of Travel & Tourism Marketing*, *34*(3), 357-368.

Payne, K. (2017). *The broken ladder: How inequality affects the way we think, live, and die*. Penguin.

Pen, J. (1971). *Income distribution: facts, theories, policies.* Praeger.

Piketty, T. (2013). *Capital in the twenty-first century.* Harvard University Press.

Piketty, T. (2015). About capital in the twenty-first century. *American Economic Review*, *105*(5), 48-53.

Piketty, T. (2018). *Brahmin Left vs Merchant Right: Rising Inequality and the Changing Structure of Political Conflict*. Mimeo, Paris School of Economics.

Piketty, T., & Saez, E. (2003). Income inequality in the United States, 1913–1998. *The Quarterly journal of economics*, *118*(1), 1-41.

Piketty, T., & Zucman, G. (2014). Capital is back: Wealth-income ratios in rich countries 1700–2010. *The Quarterly Journal of Economics*, *129*(3), 1255-1310.

Piketty, T., Saez, E., & Zucman, G. (2018). Distributional National Accounts: Methods and Estimates for the United States. *The Quarterly Journal of Economics*, *133*(2), 553–609.

Plunging response rates to household surveys worry policymakers. (2018, May) *The Economist.*

Ravallion, M. (1995). "Growth and poverty: Evidence for developing countries in the 1980s." *Economics letters* 48.3-4: 411-417.

Ravallion, M. (2018). Inequality and globalization: A review essay. *Journal of Economic Literature*, *56*(2), 620-42.

Rognlie, M. (2016). Deciphering the fall and rise in the net capital share: accumulation or scarcity?. *Brookings papers on economic activity*, 2015(1), 1-69.

Saez, E. (2004). Reported Incomes and Marginal Tax Rates, 1960-2000: Evidence and Policy Implications. In *NBER Chapters* (pp. 117–174). National Bureau of Economic Research, Inc.

Saez, E., & Zucman, G. (2016). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics*, *131*(2), 519-578.

Sen, A., Sen, M. A., Amartya, S., Foster, J. E., & Foster, J. E. (1997). *On economic inequality*. Oxford University Press.

Shannon, C. E. (1948). A note on the concept of entropy. *Bell System Tech. J*, *27*(3), 379-423.

Singer, P. (1995). *How are We to Live?: Ethics in an Age of Self-interest*. Prometheus Books.

Smith, B., & Wills, S. (2018). Left in the dark? oil and rural poverty. *Journal of the Association of Environmental and Resource Economists*, *5*(4), 865-904.

Smith-Clarke, C., Mashhadi, A., & Capra, L. (2014, April). Poverty on the cheap: Estimating poverty maps using aggregated mobile communication networks. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 511-520). ACM.

Solt, F. (2009). Standardizing the world income inequality database. *Social Science Quarterly*, *90*(2), 231-242.

Soto, V., Frias-Martinez, V., Virseda, J., & Frias-Martinez, E. (2011, July). Prediction of socioeconomic levels using cell phone records. In *International Conference on User Modeling, Adaptation, and Personalization* (pp. 377-388). Springer, Berlin, Heidelberg.

Steele, J. E., Sundsøy, P. R., Pezzulo, C., Alegana, V. A., Bird, T. J., Blumenstock, J., ... & Hadiuzzaman, K. N. (2017). Mapping poverty using mobile phone and satellite data. *Journal of The Royal Society Interface*, *14*(127), 20160690.

Stephens-Davidowitz, S. (2014). The cost of racial animus on a black candidate: Evidence using Google search data. *Journal of Public Economics*, *118*, 26-40.

Stephens-Davidowitz, S., & Pinker, S. (2017). *Everybody lies: big data, new data, and what the internet can tell Us about who we really are*. New York: HarperCollins.

Stevenson, B., & Wolfers, J. (2008). *Economic growth and subjective well-being: reassessing the Easterlin paradox* (No. w14282). National Bureau of Economic Research.

Stout, L. A. (2012). New Thinking on "Shareholder Primacy." *Accounting, Economics, and Law: A Convivium*, *2*(2), 1–22.

Suwa-Eisenmann, A., & Verdier, T. (2007). Aid and trade. *Oxford Review of Economic Policy; Oxford*, *23*(3), 481.

Tella, R. D., MacCulloch, R. J., & Oswald, A. J. (2003). The Macroeconomics of Happiness. *The Review of Economics and Statistics*, *85*(4), 809–827.

Temkin, L. S. (1986). Inequality. *Philosophy & Public Affairs*, 99-121.

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

Weidmann, N. B., & Schutte, S. (2017). Using night light emissions for the prediction of local wealth. *Journal of Peace Research*, *54*(2), 125-140.

Williams, B. (1962). The Idea Of Equality. In Philosophy, Politics, And Society, 110-131. London: Basil Blackwell.

Young, M. D. (1958). *The Rise of the Meritocracy, 1870-2033: An Essay on Education and Quality*. Thames and Hudson.

Zucman, G. (2013). The missing wealth of nations: Are Europe and the US net debtors or net creditors?. *The Quarterly journal of economics*, *128*(3), 1321-1364.

Zucman, G. (2019). Global wealth inequality (No. w25462). *National Bureau of Economic Research.*

Appendix
A1 *Traditional Data sources*

A1.1 Income

Most national statistics agencies collect income distributional data from a range of censuses, surveys and national accounts, and compute their preferred measure(s) of income inequality. There are numerous datasets that collect and harmonize these data from national statistics offices. The Luxembourg Income Study (LIS) database contains observations for 49 rich and middle-income countries. LIS estimates are calculated using a uniform set of assumptions and with harmonized microdata that maximizes comparability across countries, but this comes at the expense of global coverage. The Standardized World Income Inequality Database (SWIID), documented in Solt (2009), includes at least one Gini observation (both before and after taxes and transfers) for 192 countries since 1960, though this expanded coverage comes at the expense of reduced comparability across countries and years relative to the LIS[15]. The World Bank's PovcalNet provides several different inequality measures drawing mainly from the Living Standards and Measurement Surveys (LSMS) conducted in emerging and developing countries. PovcalNet provides distributional data for 127 countries going back to the 1980s, though the coverage over time varies widely by country.

Finally, use of tax records data to measure inequality has become increasingly common. The main advantage of tax data is that many countries have kept records going back to the early 1900s or earlier, allowing researchers to track inequality trends over very long periods. Saez, Piketty and Atkinson (2011) for example, use tax records data to track the incomes of the richest 10 to 0.01 percent of the population in 30 OECD countries, finding that the share of national income has dramatically increased for the top 1% but not for other groups.

A1.2 Wealth

In recent years, statistical agencies have begun publishing retrospective data on national stock accounts, including statistics on the market value of nonfinancial and financial assets and liabilities held by different sectors of the economy. Piketty and Zucman (2014) put together a macro-historical data set on wealth and income for the top eight developed economies in the world. The data covers the period 1970-2010 for these economies, and data goes back to 1700 for US, UK, Germany and France. They find that wealth-to-income ratios in every country have risen gradually from about 200-300% in 1970 to 400-600% in 2010, a return to the high ratios observed in Europe ion the 18th and 19th centuries. Their explanations for this trend is a long-run asset price recover and the slowdown of productivity and population growth. In other words: "capital is back because low growth is back".

The OECD has begun collecting data on the distribution of household wealth for 18 OECD countries using a host of different data sources (OECD Wealth Distribution Database). The concept of wealth used here is the 'ownership of economic capital' and excludes other types of capital (such as human, social and collective assets). Some have used national accounts data on wealth (Piketty and Zucman, 2014), while others have used data on individual tax returns (Kopczuk and Saez, 2004; Saez and

[15] Distributional income data from specific regions include the Eurostat Income Inequality Statistics, based on European Union Statistics on Income and Living Conditions (EU-SILC); the OECD's Income Distribution Database (IDD); and the Socio-Economic Database for Latin America and the Caribbean (SEDLAC).

Zucman, 2016; Piketty, 2013, Piketty and Saez, 2003). For developing countries, the Demographic and Health Surveys (DHS) Program is primarily focused on measuring health outcomes, but also collects information on several types of asset wealth and constructs a household-level asset wealth index. As of this writing it has conducted 300 representative surveys in over 90 developing countries.

A1.3 Attitude/opinion Surveys

The perception of inequality may matter just as much as its objective measurement, particularly given the social and political implications this may have. In addition, So why don't we just ask people how unequal they think society is, particularly given the various challenges in measuring inequality? Firs, there is growing evidence that people on average misperceive the actual level of inequality, typically underestimating both wealth and income inequality in their country (Hauser and Norton, 2017; Kiatpongsan and Norton, 2014). However, given the important socio-political implications, people's attitudes and opinions about inequality, regardless of their accuracy, may still be important. There is a long-standing tradition of this in the measurement of poverty, in which surveys on what people believe to be an adequate income are used to construct poverty lines (see Goedhart et al., 1977), and inequality, in which the 'economic ladder' question is often used[16]. Three challenges arise in these survey-based measures. The first is adaptation: the longer an individual is subjected to a particular condition, the more they may become adapted to life with said condition. This means that measures of poverty of inequality based on these types of surveys are biased to fall over the long-run. The second issue is how to make comparisons across societies if people have different interpretations of what inequality is. The third is a practical issue of who should collect this data, and how to get global coverage? We do not have answers to these questions, but believe that more work is needed to understand more the usefulness of opinion surveys on inequality.

[16] The 'economic ladder' question is usually of the form: "On a scale of one to ten, where one stands for the poorest level of society and ten the richest, where do you place yourself?"