Working Paper

Measuring Inequality using Geospatial Data

Author(s):
Kingeski Galimberti, Jaqueson; Pichler, Stefan; Pleninger, Regina

Publication Date:
2021-03

Permanent Link:

https://doi.org/10.3929/ethz-b-000473903 →

ETH zürich

KOF Swiss Economic Institute

Measuring Inequality using Geospatial Data

Jaqueson K. Galimberti, Stefan Pichler and Regina Pleninger

KOF Working Papers, No. 493, March 2021

KOF

KOF

ETH Zurich
KOF Swiss Economic Institute
LEE G 116
Leonhardstrasse 21
8092 Zurich, Switzerland

Phone +41 44 632 42 39
Fax +41 44 632 12 18
www.kof.ethz.ch
kof@kof.ethz.ch

Measuring Inequality using Geospatial Data**

Jaqueson K. Galimberti†[1,2,3], Stefan Pichler[2] and Regina Pleninger[2]

[1]Auckland University of Technology
[2]KOF Swiss Economic Institute, ETH Zurich
[3]Centre for Applied Macroeconomic Analysis, Australian National University

Abstract

The main challenge in studying economic inequality is limited data availability, which is particularly problematic in developing countries. We construct a measure of economic inequality for 234 countries/territories from 1992 to 2013 using satellite data on night lights and gridded population data. Key methodological innovations include the use of varying levels of data aggregation, and a calibration of the lights-prosperity relationship to match traditional inequality measures based on income data. We obtain a measure that is significantly correlated with cross-country variation in income inequality. We provide three applications of the data in the fields of health economics and international finance. Our results show that light- and income-based inequality measures lead to similar results in terms of cross-country correlations, but not for the dynamics of inequality within countries. Namely, we find that the light-based inequality measure can capture more enduring features of economic activity that are not directly captured by income.

Keywords: Nighttime lights, inequality, gridded population

JEL classification: D63, E01, I14, O11, O47, O57

**We thank Bruno Caprettini, Florian Eckert, Vera Eichenauer, Martin Karlsson, Philip Vermeulen, and Nicolas Ziebarth for helpful comments and suggestions. Moreover, we thank participants at DENS 2020 in St. Gallen, YSEM 2021 in Zurich and the research seminars at the Competent in Competition and Health center (CINCH) in Essen, at Auckland University of Technology, and at the University of Auckland. The subnational borders data used in this paper were accessed during a visit to the Center for International Earth Science Information Network (CIESIN), at Columbia University. We thank Kytt MacManus and Greg Yetman for the hospitality. We also thank the NASA Socioeconomic Data and Applications Center (SEDAC) for the support with the distribution of the data set generated by this project, which can be accessed at `https://www.ciesin.columbia.edu/data/global-geospatial-inequality/`. This project received funds from the MTEC Foundation Grant, for which we gratefully acknowledge the support.

†*Corresponding author*: Auckland University of Technology, School of Economics, B-31 Economics, Private mail bag 92006, Auckland, 1142, New Zealand. Email addresses: jaqueson.galimberti@aut.ac.nz (J. K. Galimberti), pichler@kof.ethz.ch (S. Pichler), pleninger@kof.ethz.ch (R. Pleninger).

1 Introduction

The past decades have witnessed a significant increase in economic inequality with important social and economic consequences (see, e.g., Piketty and Saez, 2014; Lakner and Milanovic, 2016). As a result, the study of inequality gained relevance within the economics profession, both in terms of studying determinants of inequality (see e.g. Milanovic, 2005; Acemoglu et al., 2015) and the implications of rising inequality (Deaton, 2003; Easterly, 2007). However, one important constraint in the study of inequality is the limited availability of consistent data and measurements at a global scale. In particular, sources and methods used to construct global databases of inequality tend to vary substantially in quality and availability across countries and within countries over time (Atkinson and Brandolini, 2001).

Traditional measures of inequality tend to rely on a mixture of sources of national accounts and household survey data. As noted in the literature, both sources are prone to design differences and scattered availability, which is especially true for survey data (Deaton, 2016). Moreover, while household surveys tend to under-sample richer households from the population (Deaton, 2005), recent advances with the incorporation of data from tax records are still subject to tax evasion and other consistency issues (Galbraith, 2019). Finally, the informal sector and shadow economy transactions pose additional threats to the reliability of these data (Alstadsæter et al., 2019). These difficulties can put the results of comparative analysis of inequality data from different sources into question, as these results could be a product of measurement errors instead of genuine effects.

This paper proposes an alternative measure to circumvent these issues by constructing an internationally comparable measure of economic inequality based on geospatial data. In particular, we use worldwide satellite data on nighttime light emission as a proxy for economic prosperity and match them with data on geo-located population counts to construct Gini-coefficients, based on light per capita. In the context of the issues discussed above, the greatest advantage of this approach is the consistent coverage provided by the geospatial

source data across countries. Our novel database covers a balanced global sample of 234 countries and territories on a yearly basis from 1992 to 2013, yielding a total of 5,148 observations. Moreover, in contrast to previous literature, our database is not skewed towards high income countries, providing equal coverage across virtually all populated regions in the world.

One potential drawback in our approach lies in the ambiguity on what is actually measured by the night lights data. Here, we argue that our measures can capture multiple aspects of inequality, including the distribution of income, consumption and wealth, as well as investment and infrastructure expenditures. Besides, in contrast to other sources of economic data, the informal activities associated with each of these aspects of inequality are less likely to be concealed from the satellite measurements. Our measure provides a combination of these aspects and thus, following the discussion of Sen (1997), reflects overall economic inequality rather than only income inequality. Given these features, our measure can be a viable alternative for researchers that are interested in studying overall economic inequality, rather than a particular aspect of inequality such as income or wealth inequality. Moreover, due to the global coverage our measure is particularly useful for researchers interested in inequality across developing countries, where no or only very few traditional inequality measures are consistently available.

Night lights data have found fruitful applications in the recent economics literature, particularly as an alternative proxy for economic activity. Typical applications based on these data include the evaluation of the accuracy of national income accounts (e.g., Henderson et al., 2012; Nordhaus and Chen, 2015; Pinkovskiy and Sala-i Martin, 2016), the geographic mapping of economic activity (e.g., Mellander et al., 2015; Bickenbach et al., 2016; Henderson et al., 2018), regional development analysis (e.g., Michalopoulos and Papaioannou, 2013), and, more recently, macroeconomic forecasting (e.g., Galimberti, 2020). In the area of inequality, several studies have proposed the use of night lights in conjunction with population data to derive measures of regional inequality in development, well-being, and income (see, e.g.,

Elvidge et al., 2012; Alesina et al., 2016; Lessmann and Seidel, 2017, and further discussion in the next section).

We complement this literature on five methodological dimensions. First, we consider different sources of population data. In particular, we combine data from the Gridded Population of the World (GPW) dataset, which is strictly based on population census data, with more granularly disaggregated population counts from the LandScan database. Second, these data allow us to construct measures at different levels of geographical aggregation, consistent with the resolution of each source. With a resolution ranging from 30-by-30 arcsecond cells (or pixels – these cover approximately one square kilometer at the equator), to whole census areas within each country, we obtain a proxy for inequality at a more granular level than the previous regional measures in the literature. Considering the substantial variation in accuracy of population location over the years across these data sources, this mixing approach also allows us to obtain more time-consistent measures.

Third, we also innovate with respect to the exploration of alternative versions of the night lights data. In particular, we adopt a version of the data that accounts for more sparsely populated regions that tend to emit less stable lights – we also consider alternatives that correct for sensor saturation at brighter locations, but find that the effects of top-coding are of secondary importance here. Fourth, in order to translate light intensities into measures of economic prosperity, we adapt a constant elasticity approach to be uniformly applied at the pixel level. Here, the relationship between lights and prosperity is regulated by an exponential scaling parameter. Finally, we develop an agnostic approach to the calibration of this relationship, particularly geared towards the measurement of income inequality.

Our approach yields several geospatial inequality measures as a product of different combinations of data sources and light intensity scaling factors that can be used in the calculation of Gini-coefficients. In order to obtain a parsimonious composite inequality measure, we combine all these geospatial Gini-coefficients by weighting them to maximize their correlation with a benchmark income inequality measure. Our approach of mixing

multiple light-based measures for a common target, economic inequality, can be related to a long standing literature on measurement error (see, e.g., Browning and Crossley, 2009), according to which several error-prone measures can provide a better measurement than a single one. Our use of income inequality as a reference for calibration is due to its greater availability as well as because income inequality is an important component in the broader definition of economic inequality. For that purpose, we adopt an established source in the literature, namely the Standardized World Income Inequality Database (SWIID), developed by Solt (2016, more details in the next section). In order to deal with the potential measurement errors discussed above, we use quality indicators from the database as weights in this estimation exercise, giving a lower weight to more uncertain datapoints.

Balancing between cross-country and within-country correlations, we obtain a composite measure that is significantly correlated with income inequality, particularly along the cross-country dimension. This is in contrast to previous literature that found light-based measures of inequality to have low correlation with income inequality (Elvidge et al., 2012). Hence, our methodological innovations allow us to obtain a clearer identification of the type of economic inequality that is measured through lights data. We also find that the light-based measures suggest generally higher levels of inequality than traditional income inequality measures. The latter finding is suggestive to our interpretation that light-based measures capture multiple aspects of inequality, such as access to public services/infrastructure and wealth, that are likely missing in income-based measures.

Finally, we use the generated inequality measure in three applications. The inequality literature has pointed to health shocks (Alam and Mahal, 2014) and financial liberalization (De Haan and Sturm, 2017) as important drivers of inequality. We revisit this literature and estimate the relationships between economic inequality and out-of-pocket health expenditure, epidemics, and financial liberalization. In particular, we compare the results for income- and light-based Gini-coefficients. We find that cross-country results between the two measures are very similar. In contrast, absorbing country-specific averages through fixed effects leads

to different results, which, we argue, can be explained by light-based measures incorporating income inequality as well as other more persistent aspects of economic activity, such as housing and infrastructure investments. Moreover, across all applications we are able to investigate a more complete sample of country-year observations using light-based inequality as compared to income-based data – e.g., for epidemics we observe an increase from 3,278 to 5,148 observations. This larger sample size is due to the extended number of countries as well as observations within countries over time. Finally, we generally find that estimates with light-based Gini-coefficients are robust across the different samples, as well as across alternative light weighting schemes.

The remainder of this paper proceeds as follows. The next section gives a short overview of the related literature. In Section 3, we provide a description of the data and adjustments needed to construct the light-based inequality measure in Section 4. Three applications of the data are provided in Section 5. The last section concludes.

2 Related Literature

Previous studies on inequality rely on different sources of data that range from household-level surveys to aggregate regional and country-year databases. In this section, we discuss common sources of data on income inequality and their (un)availability issues. We then turn to a review of studies more directly related to our approach of using geospatial data for economic measurement.

2.1 Income inequality data

The most common sources of data on income inequality are the World Income Inequality Database (WIID), from UNU-WIDER, and the Standardized World Income Inequality

Database (SWIID).[1] The WIID combines data from the OECD, Eurostat, Luxembourg Income Study (LIS), World Bank and household surveys. It contains distributional data, such as Gini-coefficients, percentiles and other income information, for 179 countries up to 2015. However, the database is very unbalanced: in the 1992-2013 sample period, missing country-year points account for about 54% of the total data available for countries included in the database.

The SWIID database is developed and frequently updated by Solt (2016) with the aim of improving cross-country coverage while keeping the comparability between countries and years. The sources are the WIID database, LIS, other cross-national and national databases as well as data from scholarly articles. The SWIID database provides a significant improvement in coverage relative to the WIID, with Gini-coefficients for a total of 191 countries over the period from 1960 to 2016. Nevertheless, the percentage of missing values for the 1992-2013 years still amounts to about 20% of a balanced sample. We will use the SWIID data as a baseline reference for the calibration of our own new measure and for comparative purposes in the applications.

Solt (2016) uses two LIS series as baselines for standardizing the source data by generating model-based multiple imputation estimates for missing observations. The standard deviation over these imputed values can be used to capture the quality of the data, where more precise estimates are related to less dispersed imputations. Jenkins (2015) gives an overall assessment of the WIID and SWIID data, pointing to caveats about the SWIID's imputation approach. In particular, a potential misspecification of the imputation model can bias estimates of inequality and lead to uncertainty about what is actually being measured. However, Jenkins (2015) also finds that constructing averages across the imputed values leads to similar standard errors as in the case where the imputation variability is taken into account. This procedure reduces potential biases and is, thus, applied in our subsequent analysis. Finally,

[1]See Ferreira et al. (2015) for a comprehensive review of cross-national databases of income inequality.

with the goal of analyzing the influence of the imputation model on our results, we use the WIID data as a baseline reference in a robustness check. We provide more details on this in Section 4 below. Our findings show a high correlation between WIID- and SWIID-based measures, hence we conclude the (potential) bias from the imputation model has a negligible effect on our measure.

Another important difference of the SWIID database is that it provides distinct inequality measures for market-based and after-tax income. Market-based Gini-coefficients are calculated on pre-tax income, which excludes any form of redistribution. In contrast, net income inequality is based on income after taxes and social benefits. The sample size and the distinction between market- and net-based inequality measures constitute important advantages of the SWIID compared to the WIID data. For our purposes, we will adopt the net income measure because it reflects income available to individuals and is therefore more relevant for light emission.[2]

2.2 Night lights data in economics

Seminal studies on the use of night lights data in economics have focused on the data's capabilities as a proxy for economic activity, especially in the context of national statistics. The contribution of night lights data as an alternative to traditional national accounts data was, among others, analyzed by Chen and Nordhaus (2011). Their goal was to determine whether night lights add value to the conventional national accounts statistics. Similarly, Henderson et al. (2012) proposed a statistical framework to combine night light measurements with conventional national income statistics and generate an improved estimate of economic growth. Their overall finding is that night lights data improve data availability and quality, especially in countries with less developed statistical agencies. Nordhaus and Chen (2015)

[2]In another robustness check, we considered using the market-based inequality measures as reference and found that the correlations with the implied lights-based measures are less than half the magnitude of the correlations obtained using the net income measures as reference. These results are available upon request.

go one step further and argue that lights data can offer additional information for a more granular analysis of regional output within developing countries. More generally, see also Donaldson and Storeygard (2016) for a review of applications of remotely sensed data in economics.

In the context of inequality, there have been several studies proposing the use of the night lights data to construct alternative measures of inequality. One earlier contribution is provided by Elvidge et al. (2012), who combine night lights with gridded population data to construct a light-based Gini-index of development. They find that the resulting measure of inequality tends to be inversely related to different measures of development, but poorly correlated with a traditional income-based Gini-coefficient.

In a related paper, Mveyange (2015) looks at regional inequality in Africa for 1,235 regions by relying on light and population data. In contrast to Elvidge et al. (2012), Mveyange finds a strong correlation between light inequality and income inequality calculated at the regional level. Following a similar methodology, Lessmann and Seidel (2017) use night lights and population data to construct a panel dataset of regional incomes worldwide, covering 180 countries from 1992 to 2012. They then use the new measure to study cross-country convergence in the dispersion of regional incomes over time, finding evidence of an N-shaped relationship between development and regional inequality.

Going beyond the scope of income inequality, Alesina et al. (2016) discuss the use of night lights for constructing a measure of ethnic inequality. For this, they combine linguistic maps on the spatial distribution of ethnic groups within countries with corresponding measures of night lights density. Taking advantage of the regional differences in the luminosity data, the authors argue that their inequality measure captures differences in well-being across different ethnicities, particularly through varying levels of public goods provision. They conclude that differences in geographic endowments can explain a large fraction of the variation in economic inequality across groups. Related to this finding, Hodler and Raschky (2014) observe that

birth regions of incumbent political leaders have more intense nighttime lights, which is likely to affect regional inequality.

Another related contribution is provided by Weidmann and Schutte (2017), who use night light emissions to predict local wealth at the level of villages or neighborhoods across a sample of 39 developing countries. They find that light emissions provide highly accurate predictions of wealth measurements derived from geo-referenced survey data. Although these authors do not calculate measures of inequality, it is interesting to note that income and wealth inequality are correlated, though the level of wealth inequality tends to be considerably higher across countries (see Davies et al., 2009; Piketty and Saez, 2014).

Our paper complements this literature in several dimensions. First, we consider alternative sources of data, particularly using a mixture of data on gridded population. As described in the next section, our approach takes advantage of cross-source variation in the precision of geolocated population data, both across space and over time. Second, we also explore varying levels of geographical aggregation in the construction of our inequality measures. Another key innovation is that we consider a version of the night lights data that provides a better account for sparsely populated regions that emit less stable lights. Finally, our measures are based on an agnostic calibration of the relationship between lights and economic activity, which, in turn, is informed by a comparison to the income inequality measure discussed above. Taking all of these innovations together, we obtain a correlation of approximately 0.5 with SWIID. This is a much higher correlation as compared to the earlier literature. Alesina et al. (2016), for instance, find a correlation of 0.25-0.3 between their ethnic inequality measure and the SWIID. Correlation estimates in Mveyange and Lessmann and Seidel (2017) between light- and income-based inequality are in a similar range. We describe these methodological innovations in the following sections.

3 Source Data and Measurement Issues

In order to construct geospatial measures of inequality, we rely on nighttime lights and population data. These data come from different sources and in varying formats, naturally ensuing different issues on their use for inequality measurement. In this section, we describe these source data and relevant measurement issues.

3.1 Night lights data

Source and construction The night lights data are obtained from the Earth Observation Group (EOG) at the National Oceanic and Atmospheric Administration's (NOAA) National Centers for Environmental Information (NCEI). We use the data based on the Operational Linescan System (OLS) instruments, which have been designed to capture images of cloud cover by the United States Air Force Defense Meteorological Satellite Program. These data are available at the annual frequency between 1992 and 2013 with a full coverage of Earth's surface along the longitudinal dimension between 75 degrees north and 65 degrees south in latitude. Figure 1 illustrates the visual features of these data over selected regions.[3]

The satellite data were processed by EOG scientists into the form of global composite images representing the average intensity of lights captured by the OLS sensors during nighttime (around 9pm local time) over the year. The intensity of the night lights radiance are converted into 6-bit digital number (DN) values, ranging between 0 and 63, and allocated over a global grid of 30 arc second cells (or pixels) according to their geographic location. The annual measurements are based on averages of cloud-free observations, also discarding images affected by sunlight, moonlight, and aurora lighting (Elvidge et al., 2003). Hence, the

[3]The OLS-based night lights data have been discontinued after 2013, being replaced by the more recently launched Visible Infrared Imaging Radiometer Suite (VIIRS) sensors (see Miller et al., 2013). At the time of writing, there were still open issues with respect to how VIIRS compares to OLS when merging the two datasets. Hence, we decided to focus only on OLS data, also because these are the longest available series of lights data covering the period prior to 2013, leaving an extension of our methodology to VIIRS data for future research.

Figure 1: Illustrative Extracts of Average Night Lights

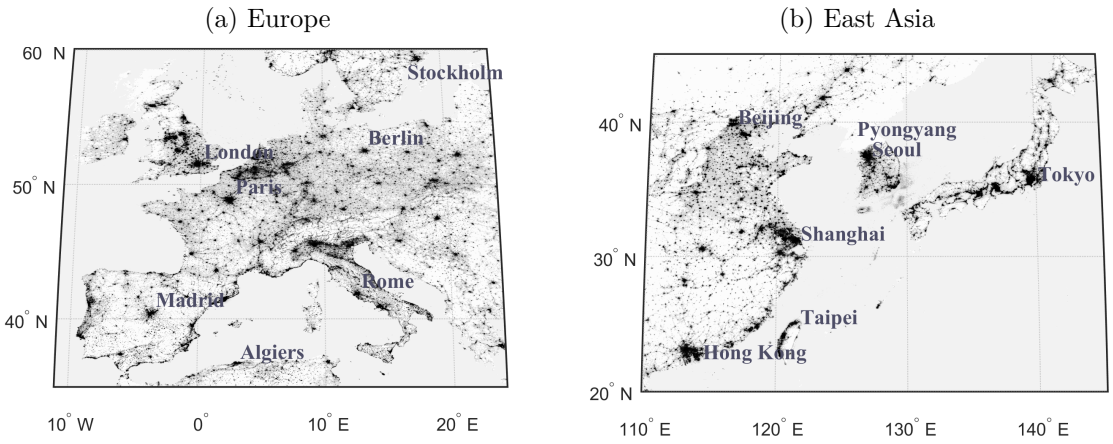

Notes: Based on 2013 data from the Operational Linescan System (OLS) obtained from NOAA-NCEI. The darker the pixel, the higher the light intensity.

number of underlying observations in the annual lights measurements varies across countries, particularly with respect to their latitude – more sunlight and aurora effects in Nordic countries – and the prevalence of areas covered by rainforest – more clouds over tropical countries. The statistics in the first row of Table 1 provide a summary of the cross-country variation in this variable.

Stable lights and low-coding There are two main versions of the OLS data, one averaging the raw visible band DN values, while the other, so-called stable lights, is produced after filtering out non-persistent and "low-coded" light intensities – these are mostly locations with light intensity between DN=1 to DN=3. The latter is the version that is mostly used in social sciences applications, because it provides a clearer identification of lights from urban centers (Elvidge et al., 2003). Nevertheless, as the statistics in the second and third rows of Table 1 indicate, the stable lights algorithm can have a strong effect on the measurement of low-coded signals, particularly over the less densely populated regions. In fact, these statistics reveal that a substantial share of the population for some countries in our sample is located in these regions. Assuming that these locations have zero lights, as implied by the stable lights data,

Table 1: Cross-Country Data Summary Statistics

	Min	Quartiles			Max
		Q25	Q50	Q75	
Avg. cloud-free obs. per pixel (OLS)	6.42 Greenland	39.53 Colombia	50.18 Norfolk Isl.	58.03 Qatar	74.84 Mauritius
Fraction unlit pixels (OLS stable lights DN=0)	0 Singapore†	0.27 Switzerland	0.75 Norway	0.95 Gambia	1.00 Tokelau
Pop. share in unlit pixels (LSC pop., OLS stable lights DN=0)	0 Singapore†	0.03 Portugal	0.19 Colombia	0.57 Gabon	1.00 Tokelau
Fraction top-coded pixels (OLS DN=63)	0 Liberia†	0 Fiji†	0.0002 Macedonia	0.0014 Tunisia	0.578 Singapore
Pop. share in top-coded pixels (LSC pop., OLS DN=63)	0 Guinea†	0 Somalia†	0.040 Namibia	0.163 Russia	0.922 Bahrain
N. census areas (GPW)	2 Norfolk Isl.†	29 Botswana	153 Nicaragua	739 Nigeria	10.5M U.S.A.
N. populated pixels (LSC)	9 Monaco	3032 W. Samoa	78,047 Oman	368,626 Cote d'Ivoire	9.1M China

Notes: All statistics are based on country averages over time; † indicates a selected country among others showing an equal value for the corresponding statistic.

would distort the measurement of inequality for such countries. Therefore, we adopt the average visible lights version of the OLS data.[4]

Sensor saturation and top-coding Another issue affecting the scale of the OLS data is sensor saturation over brighter sources of lights: signals that exceed the sensor's detection range are recorded with the highest DN value in the OLS scale (i.e., 63), which is referred to as "top-coding" (Hsu et al., 2015). As the statistics in the fourth and fifth rows of Table 1 indicate, this issue can be a concern for highly urbanized countries, especially developed small island countries such as Bahrain and Singapore. By its nature, this problem is of a lesser concern to less developed countries, with over a quarter of the countries in our sample not showing any top-coded pixels during our sampled period. In fact, adopting adjustments proposed in the literature (see Bluhm and Krause, 2018) we find that top-coding has only

[4]We also calculated inequality measures on the basis of the stable lights data. While the overall correlation between these two alternative measures is high ($\rho \in [0.75 - 0.85]$), there are important differences mainly for low income countries. These alternative measures are available upon request.

minor effects on our inequality measures.[5] Hence, we continue to adopt the uncorrected original OLS data.[6]

Lights and economic activity The underlying assumption in the use of the night lights data for economic measurement is that locations with brighter lights are likely more economically prosperous than dimmer ones. This relationship, nevertheless, is not necessarily linear in levels and across scales of geographical aggregation. Varying industrial compositions, rates of urbanization, and cultural uses of light can affect how observed lights translate into economic development across regions. In fact, previous research suggests that, at the national aggregate level, a (nonlinear) constant elasticity relationship between lights and income is more appropriate (Henderson et al., 2012). Moreover, there is mounting evidence that such a relationship can be heterogeneous across countries, depending on levels of development and the structure of the economy (Hu and Yao, 2019; Galimberti, 2020). The relationship between lights and economic activity is also notably more uncertain at higher levels of disaggregation (Nordhaus and Chen, 2015; Bickenbach et al., 2016).

In this paper, we adopt a constant elasticity approach to process the light intensities into measurements of economic prosperity at the pixel level. Namely, we assume a nonlinear relationship of the form

$$x_i = DN_i^{\kappa}, \tag{1}$$

[5]Particularly, we find that: (i) the overall correlation between inequality measures based on non-corrected versus corrected lights data is equal to 0.89; (ii) the overall difference between these two versions, averaged across all countries in our sample, is close to zero.

[6]Three additional issues known to affect the OLS lights measurement are: (i) gas flares, (ii) the lack of an inflight calibration mechanism across satellites, and (iii) over-glow/blurring (see Gibson et al., 2020). On the first two, we follow the literature by excluding areas identified with gas flares from our calculations, and using an invariant region method to intercalibrate the data across the different satellites and years (see Elvidge et al., 2009, for further details). Regarding over-glow/blurring, an optical problem that causes lights to be wrongly attributed to neighboring areas, we experimented with a smoothing algorithm similar to that on-board the satellites (as described by Elvidge et al., 2004) and also calculated inequality measures on the basis of data – lights and population – averaged in 5×5 blocks of pixels to make the spatial resolution between these data more consistent; the resulting measures are virtually equal to using the original data, and are available upon request.

where each pixel i's digital number (DN) on light intensity is converted to a light-based measure of economic activity, regulated by an exponential scaling parameter, κ. As reviewed above, estimation of κ at such a disaggregated level is very difficult due to the unavailability of consistent data on economic activity, x_i, as well as the likely existence of heterogeneous relationships between lights and economic activity across different locations. As noted by Henderson et al. (2012), even at the national aggregate level, estimates of κ will tend to be affected by national statistics and satellite measurement errors. In this context of ambiguity, we follow an agnostic approach to the determination of the night lights scaling factor κ, and consider multiple calibrations for this parameter. Namely, we calculate separate inequality measures for each $\kappa = \{0.5, 1.0, 1.5, 2.0, 3.0, 5.0\}$, while allowing these different calibrations to mix in a latter stage when we construct composite inequality measures.[7]

3.2 Population data

Sources and construction We consider two sources of population data. The first is the Gridded Population of the World (GPW) dataset, produced by the Center for International Earth Science Information Network (CIESIN, 2016). The GPW is constructed on the basis of population census data, collected from hundreds of organizations that include national statistics offices and other mapping agencies, matched to spatially-explicit administrative boundary data (Doxsey-Whitfield et al., 2015). The second source is the LandScan (LSC) database, produced by the Oak Ridge National Laboratory (Bright et al., 2017). In contrast to the GPW, the LandScan data are based on a multi-variable mapping approach that disaggregates census counts within administrative boundaries with the support of ancillary data, such as land cover, roads, slope, urban areas, village locations, and high resolution

[7]Remember that DN can take values between 0 and 63. For the brightest pixels, using a scaling parameter of 5 obviously translates into a very large number in terms of economic activity. However, as described more in detail below, we divide this economic activity by the population in a pixel in order to calculate the economic activity per capita. For the brightest pixels, a typical population per pixel is between 100 and 150K. Therefore, the resulting economic activity per capita used for the Gini-coefficient calculation is only around 10K.

imagery (see Dobson et al., 2000).[8] Figures 2 (b) and 2 (c) present a comparative illustration of these two data sources for the population of Gabon in 2010.

Census-level aggregation Similarly to the night lights data, the GPW and LSC population counts are distributed over 30 arc seconds grids of cells across Earth's surface. However, one important difference between the two population datasets is that the GPW is based on an areal-weighting method that uniformly distributes population counts across the gridded cells within the census area. This feature is clearly illustrated for the case of Gabon's second largest city, as depicted in Figure 2 (b): the GPW data "spread" Port-Gentil's population across the Bendje department, which is the census area covering that city. Hence, the GPW data at the cell level tend to underestimate population in more populated cells, and vice versa. Naturally, an inequality measure based on this information is likely to be biased, and for that reason we use data aggregated at the census area level for the GPW data. One potential issue with this approach is that GPW-based measures of inequality will lack a consistent definition regarding their underlying level of geographical aggregation. Namely, the geographical definition of census regions can vary in size and shape, both within and across countries. Besides, the number of census regions can also vary substantially across countries – see the statistics in the sixth row of Table 1.

Pixel-level aggregation In this context, the LSC dataset provides an interesting alternative to construct inequality measures with a more controlled degree of geographical aggregation. Relative to the GPW, the LSC stands in the opposite end of the spectrum of granular precision, as the multi-variable mapping approach provides more reliable estimates of population counts at the pixel level. Moreover, the use of pixel-level data greatly increases the number of groups we are able to use in the construction of our inequality measures; as

[8] Night lights were also among the LandScan data sources between 1999 and 2004, which can introduce biases in our inequality measures. After considering alternative extrapolation approaches, not including these years, we obtained similar results. Hence, we decided to use all available data from this source as originally provided.

Figure 2: Gabon – 2010

(a) Night Lights including Gas Flares

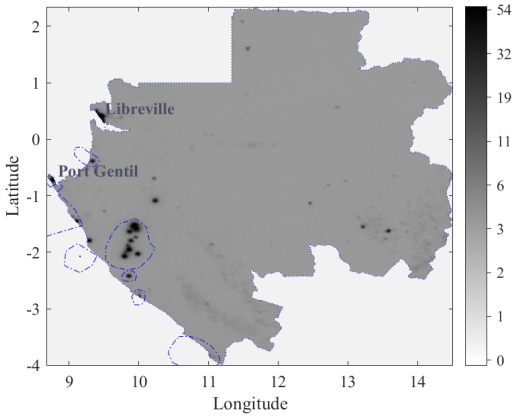

(b) GPW Population Counts

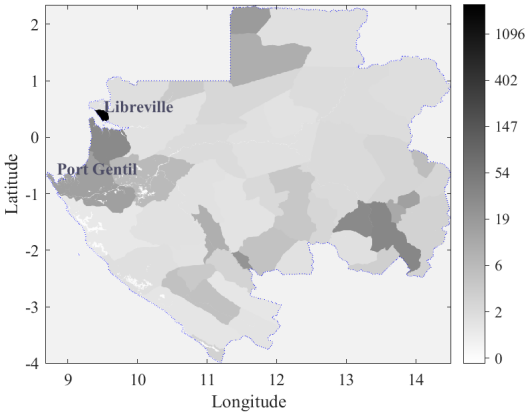

(c) LSC Population Counts

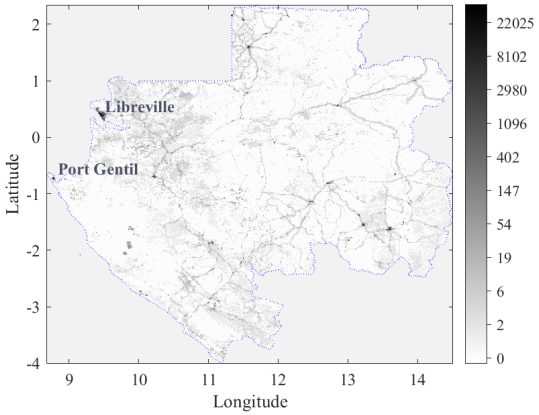

(d) Lorenz Curves

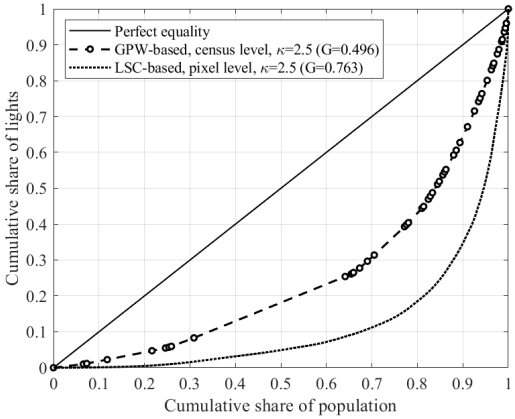

Notes: The mapped data are depicted in logarithmic scale. The blue dashed lines in panel (a) indicate areas with identified gas flaring.

the statistics in the last row of Table 1 indicate, the number of usable pixels for the median country in our sample is about 510 times larger than the corresponding median number of census areas. However, another important difference between the GPW and LSC datasets is that the latter aims to measure daytime ambient population, in contrast to GPW's focus on nighttime resident population. Hence, the greater precision of LSC comes at a cost of less consistency with the lights measurements in terms of their representative location of population.

Availability and revisions The population data also differ with respect to availability and revisions over time. The GPW is mainly updated on a decadal basis following each round of the United Nations World Population and Housing censuses.[9] We use GPW version 4.10, which is based on data from censuses occurring between 2005 and 2014, and extrapolated to produce population estimates for the years 2000, 2005, 2010, 2015, and 2020. The LSC dataset, in contrast, is available starting from the year 2000 and is updated on an annual basis, mostly to incorporate new spatial data and imagery analysis. Hence, whereas advances in high resolution imagery over the years allowed the LSC to obtain more precise measurements of population location, comparability of LSC data on a cell to cell basis across different versions is compromised. For our purposes, this means the dynamics of inequality measures derived on the basis of these data should be interpreted with caution as their underlying granular precision is changing over the years. Notice the GPW is not affected by this issue, though its dynamics are again derived only at the census level.[10]

3.3 Borders data

The geospatial data described above come in the form of global composite images. In order to construct our inequality measures, we need to extract these data for each country according to their borders. For that purpose, we adopt the border definitions from the GPW companion National Identifier Grid dataset, which contains definitions for 241 countries/territories. Moreover, for the case of inequality measures aggregated at the census level, we also need borders data at the census level. In order to maintain consistency with the construction of the GPW data, we again adopt the census borders associated with that dataset.[11] After

[9] CIESIN updates the data more frequently in the form of new release versions. However, the main versions are based on each decadal census round.

[10] In order to construct inequality measures for the broader sample of data available on night lights, we extrapolate the population data to cover an annual sample from 1992 to 2013. For that purpose, we follow a similar approach to the one used in the extrapolation of GPW population counts. See Appendix A for details.

[11] Due to license restrictions, the subnational borders data are not publicly available for some countries. We obtained these data through an on-site visit to CIESIN.

restricting our sample to countries for which there is more than one census area available, we are left with a sample of 234 countries/territories.[12]

4 Geospatial Inequality Measures

One important advantage of gridded night lights and population data is their availability at a global scale. We now attempt to capitalize on this greater availability and the properties of these data to construct alternative measures of inequality across a sample of 234 countries/territories over the period from 1992 to 2013, thereby obtaining 5,148 datapoints. The main distinguishing feature of these measures is their level of geographical aggregation, which is determined by the type of population data used to assess the spatial distribution of lights. Namely, we construct two types of measures, one aggregated at a subnational level, depending on each country's definitions of census areas, and another at a deeper level of granularity obtained with the data geo-located over grids of 30 arc second cells. We refer to these two sets of measures as geospatial Gini-coefficients. We then merge these geospatial Gini-coefficients into a single composite measure, namely the light-based Gini-coefficient, designed to proxy income inequality at the country level. In this section, we describe our approach for the construction of these geospatial inequality measures, and discuss some key descriptive statistics.

4.1 Geospatial Gini-coefficients

We measure inequality with geospatial Gini-coefficients, which are calculated on the basis of the relative distribution of lights across geographically grouped population within countries. The Gini-coefficient is a measure of dispersion commonly used to measure inequality in the distribution of income or wealth across people or groups of people. Its graphical representation

[12]The following territories had only one census area and are thus not in our sample: Anguilla, Gibraltar, Pitcairn, Saint-Barthelemy, Saint-Martin (French part), Svalbard and Jan Mayen Islands, and the Vatican/Holy See.

is given by the Lorenz curve, illustrated in Figure 2 (d), which depicts how the cumulative shares of (transformed) lights, ranked by (transformed) lights per capita (see Equation 1), evolve relative to the cumulative shares of population in a country; the Gini-coefficient is defined as the relative area between the line of perfect equality and the Lorenz curve. Mathematically, we calculate Gini-coefficients using a formulation for grouped data (see Appendix B for the derivation):

$$G = \frac{2}{N} \frac{\sum_{j=1}^{M} \left(\frac{n_j+1}{2} + a_{j-1} \right) x_j}{\sum_{j=1}^{M} x_j} - \frac{N+1}{N}, \tag{2}$$

where M denotes the total number of groups (pixels or census areas) in the country grid, N the total number of people, n_j the population in group j, and x_j the (transformed) light intensity in group j, with the index ranked in ascending order of the groups' (transformed) light intensity per capita. a_j is the cumulative sum of population up to group j, i.e., $a_j = \sum_{k=1}^{j} n_k$ for $j > 0$, and $a_0 = 0$.[13]

We calculate two main versions of geospatial Gini- coefficients by varying the level of geographical aggregation according to the source of population data. Namely, we use GPW population data to construct country-level Gini-coefficients based on data grouped at the census level, and LSC data for country-level Gini-coefficients calculated at the pixel level. By construction, these country-level Gini-coefficients capture different degrees of inequality in the spatial distribution of lights among the people that live in a country. In other terms, the geospatial Gini-coefficients capture varying degrees of the "between" versus "within" regions components of nationwide personal inequality, depending on their level of aggregation: the greater the aggregation, the greater the dominance of the regional component of inequality. In fact, notice that the decomposition of Gini-coefficients based on spatially grouped units is also affected by potential overlaps between the regional distributions of lights (Shorrocks

[13]The calculation of the Gini-coefficient with discrete data is known to be downward biased in small samples (Deltas, 2003). See Appendix B for further discussion and possible corrections. We find that such corrections have no impact on our results, hence, we continue adopting the version given by Equation (2).

and Wan, 2005). Hence, the census-level Gini-coefficients will tend to generate lower levels of inequality, due to the missing within-region inequality, while the pixel-level Gini-coefficients will tend to be more strongly affected by the measurement errors stemming from the source data, as described above. The former prediction is clearly illustrated by the comparison of these two versions of Gini-coefficients for the case of Gabon in Figure 2 (d). Summary statistics covering the whole sample of these geospatial Gini-coefficients, presented in Table C1 in the Appendix, further corroborate these arguments: (i) the census-level measures are on average lower than the pixel-level ones; (ii) the LSC-based Gini-coefficients show greater variation over time than those based on GPW.

4.2 Weighted light-based inequality measure

Other than the two versions discussed above, we also calculate geospatial Gini-coefficients for six calibrations of the lights scaling factor, κ, leaving us with a total of 12 geospatial measures of inequality. In order to synthesize these measures into a single one, we now derive a weighting scheme designed to maximize the correlation of a composite measure with a benchmark measure of inequality. In particular, we are interested in a measure that resembles measurements of income inequality, where we take the SWIID after-tax income Gini-coefficients as a reference (Solt, 2016).[14]

Weights estimation By the two-dimensional nature of our data, we can distinguish between two different correlation goals: cross-country and within-country (time-series) correlations. The cross-country correlations are calculated by matching the inequality measures on a yearly basis, while the within-country correlations are calculated by pairing the measures at the country level. We then average these correlations and estimate weights

[14]In a robustness check described below, we use the WIID as a reference (UNU-WIDER, 2017).

on our geospatial Gini-coefficients by solving the following numerical optimization problem:

$$\max_{\{\omega_i\}_{i=1}^{12}} \left\{ \lambda \overline{\rho^{cross}_{G^w,G^y}} + (1-\lambda)\, \overline{\rho^{within}_{G^w,G^y}} \right\} \tag{3}$$
$$s.t. \sum_{i=1}^{12} \omega_i = 1,\ 0 \leq \omega_i \leq 1\ \forall i,$$

where

$$G^w = \sum_{i=1}^{12} \omega_i G_i^l \tag{4}$$

is the resulting weighted light-based Gini-coefficient, obtained by weighting the 12 geospatial Gini-coefficients, denoted by G_i^l. $\overline{\rho^{cross}_{G^w,G^y}}$ and $\overline{\rho^{within}_{G^w,G^y}}$ stand for the average cross and within-country correlations, respectively, between the weighted Gini-coefficient and the SWIID income-based Gini-coefficient, G^y. Importantly, we calculate weighted correlations in order to account for the varying accuracy in the SWIID imputations, using the inverse of the SWIID estimated standard errors as weights. Finally, λ is a parameter regulating the relative relevance of the two correlation goals in the construction of the weighted measure. We solve this maximization problem for 3,287 observations, where we have observations on the SWIID income-based Gini-coefficient.[15]

Correlations trade-off Numerical estimation of (3) indicates the existence of a trade-off between how close our weighted measure can resemble the cross-country dispersion of income-based Gini-coefficients versus their correlation in terms of within-country dynamics. This trade-off is illustrated in Figure 3 (a), which also depicts the correlations between the original geospatial and the income-based Gini-coefficients. There is a clear dominance of the weighted Gini-coefficient over the original ones. Besides, the correlations between the

[15] In an earlier version of this paper, we formed country groups based on (i) country size, (ii) population density, and (iii) light density and calculated group-specific weights based on the equation above. The results using these different groups were very similar. However, since the group formation was quite arbitrary, we decided to abandon this approach. Relatedly, we experimented with estimating country-specific weights, but countries with very few or no income inequality observations proved to be significant hurdles for that approach.

Figure 3: Geospatial and Weighted Light-based Gini-coefficients

(a) Correlations with Income Gini

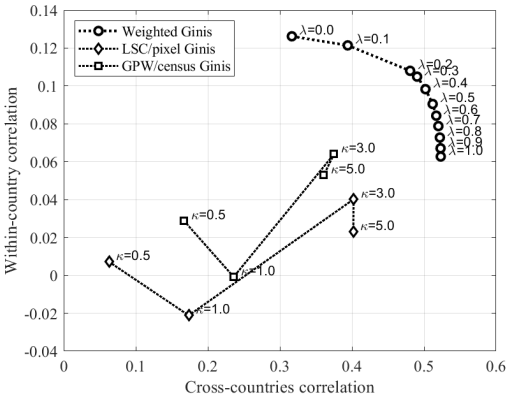

(b) Estimated Weights

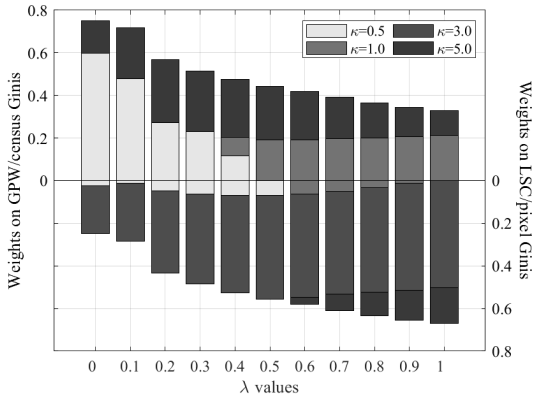

Notes: The cross and within correlations reported in Panel (a) are averaged across years and countries, respectively. In Panel (b), the top bars reflect weights on GPW/census Gini-coefficients (left scale), while the bottom bars refer to weights on LSC/pixel Gini-coefficients (inverted right scale).

weighted and the income-based Gini-coefficients are affected by λ, as expected: the cross-country correlations increase with λ, while the within-country correlations decrease with λ. Relative to previous literature, these correlations are on a higher ground. Elvidge et al. (2012), for example, find a cross-country correlation of only about 0.3 between their night light development index and income inequality. Although the choice of λ can be of secondary importance for applied purposes, the results reported in Figure 3 (a) suggest a calibration in the middle range provides an interesting compromise between cross and within correlation with the measure of income inequality. Hence, in what follows for our applications in the next section we adopt $\lambda = 0.5$ for our main results and discussion, while reporting the robustness of these results to alternative specifications in the Appendix.

Figure 3 (b) presents the estimated weights given to the original geospatial Gini-coefficients – none of the measures based on $\kappa = \{1.5, 2.0\}$ received a positive weight, hence they are not presented. It is interesting to note that the census-level GPW-based Gini-coefficients tend to receive a greater weight when targeting maximal within correlation ($\lambda = 0$), which is consistent with our analysis above indicating that this source provides

more reliable measurements over time. Nevertheless, we find that our measures capture the cross-country dispersion of income inequality more accurately, while providing less reliable proxies for the within-country dynamics in income inequality identified by the SWIID data. In particular, notice that we can easily achieve an average cross-country correlation of over 0.5, while the maximal average within-country correlation is below 0.13 – we also find a greater dispersion in the within-correlations across countries (standard deviation of about 0.5) than in the cross-correlations over the years (standard deviation of about 0.04). Further summary statistics on the weighted and income-based Gini-coefficients are presented in Appendix C.

Country-specific factors The cross-country variation in the different measures of inequality can be related to country-specific characteristics. Here, we highlight two factors that are directly related to the construction of our measures: the countries area sizes and lights density levels–although we do not pursue an exhaustive analysis of the many potentially correlated factors at this stage, we do provide additional applications on that direction in the next section. First, the relation between our inequality measures and country size is illustrated in Figure 4 (a). The light-based Gini-coefficients tend to be lower for smaller countries, while the income-based measures are virtually insensitive to country size. One potential explanation for this result, as previously alleged, is that our measures capture a broader definition of prosperity than income, which can, in turn, be related to country size. For example, provision of basic public goods such as electrification, sanitation, and transportation, can require more complex logistics in bigger countries, hence posing a greater challenge for an equitative distribution of such services across regions. Our inequality measures are consistent with this hypothesis as we observe higher levels of inequality in bigger countries. Nevertheless, for the purpose of measuring income inequality, the disparate effect of country area in our measures should, ideally, be controlled for in the applications.

Second, Figure 4 (b) shows that there is a negative relationship between measured inequality and the total lights per area observed within countries. The latter statistic

Figure 4: Cross-country Variation in Measured Inequality by Factors

(a) Area (b) Lights/Area

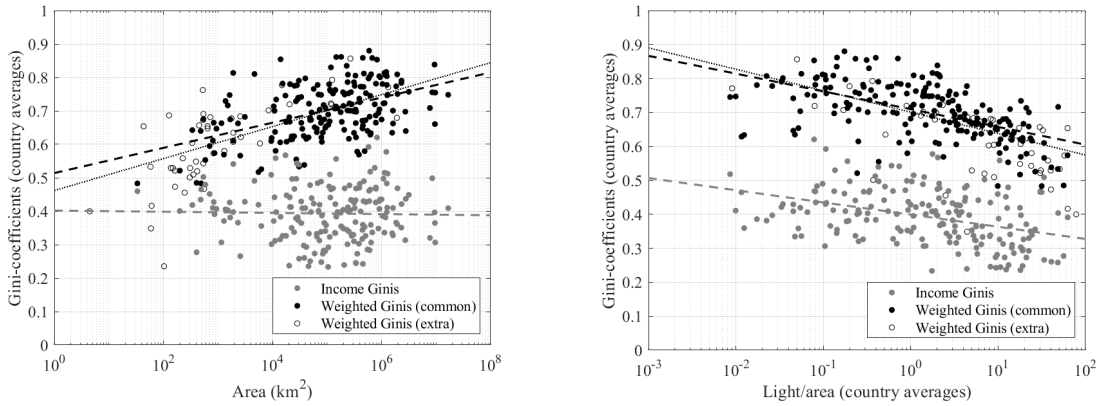

Notes: The figures show the relationship between income- and light-based Gini-coefficients and country areas (a) as well as lights per area (b). Each datapoint represents a country-level average. The weighted Gini-coefficients are obtained with $\lambda = 0.5$. Black dots denote countries that have income- and light-based Gini-coefficients and, hence, constitute a common sample. Whereas white dots correspond to countries with no income-based data. The dashed regression lines are obtained under common samples for the income and weighted Gini-coefficients, while the thin dotted lines reflect the extended sample of weighted measures.

captures a mixture of level of development and population density, as the countries in the higher end of the lights per area range include mostly European countries and small island nations/territories. Interestingly, measured inequality tends to be lower for those countries, both in terms of our weighted light-based measure and the traditional measure of income inequality. Hence, while the lights density statistic naturally carries the effects of country area pointed out above, as a proxy for development it is reassuring to find that our measures are also consistent with the prevailing view that there is a negative relationship between inequality and level of development (see, e.g., Barro, 2000).

Extended sample The estimation of the weights used to calculate our weighted inequality measure is based on a restricted sample of common observations between the SWIID and our geospatial Gini-coefficients. Namely, the estimation sample contains 3,266 observations for a total of 175 countries/territories, unbalanced over the years between 1992 and 2013. However, because these weights are assumed to be equal across countries and years, we can use our data

to extrapolate our weighted light-based Gini-coefficient to a broader, and balanced, sample of 5,148 country-year observations, covering a total of 234 countries/territories over the period from 1992 to 2013.[16]

Figure 5 illustrates the relationship between these measures of inequality, also including datapoints for which only light-based measurements are available. There is a clear difference in levels between the light- and the income-based Gini-coefficients, the former being on average about 0.3 points higher than the latter. Although these absolute level differences are less relevant for comparative analysis of inequality across countries and over time, they can account for how the lights data complement income measurements as a proxy for multiple aspects of development. Namely, the night lights data can capture varying levels of human development with respect to the geographical distribution of economic activity, infrastructure and wealth (see, e.g., Michalopoulos and Papaioannou, 2013; Henderson et al., 2018).

The distribution of the datapoints in Figure 5 show a clear positive correlation between income- and light-based Gini-coefficients. Outliers, such as Barbados, have smaller marker sizes, which indicates a lower quality of their corresponding income-based measures of inequality. In general, countries with higher income-based data quality are closer to the regression line. Altogether, these features represent evidence that our measure is closely related to income inequality, while also capturing additional aspects of economic inequality. We now turn to an assessment of these new measures by providing several applications of the data.

Robustness In order to see whether the imputation model used in Solt (2016) affect our results, we repeat the optimization problem described in equation (3), but use WIID as a reference (UNU-WIDER, 2017). This drastically reduces the number of observations to 1,777 (compared to 3,287 for the SWIID). However, the newly obtained WIID-based Gini-coefficient shows a very high correlation with the SWIID-based Gini-coefficient. For instance,

[16] A complete list of countries and territories is provided in Table D5 in the Appendix.

Figure 5: Light and Income-based Gini-Coefficients

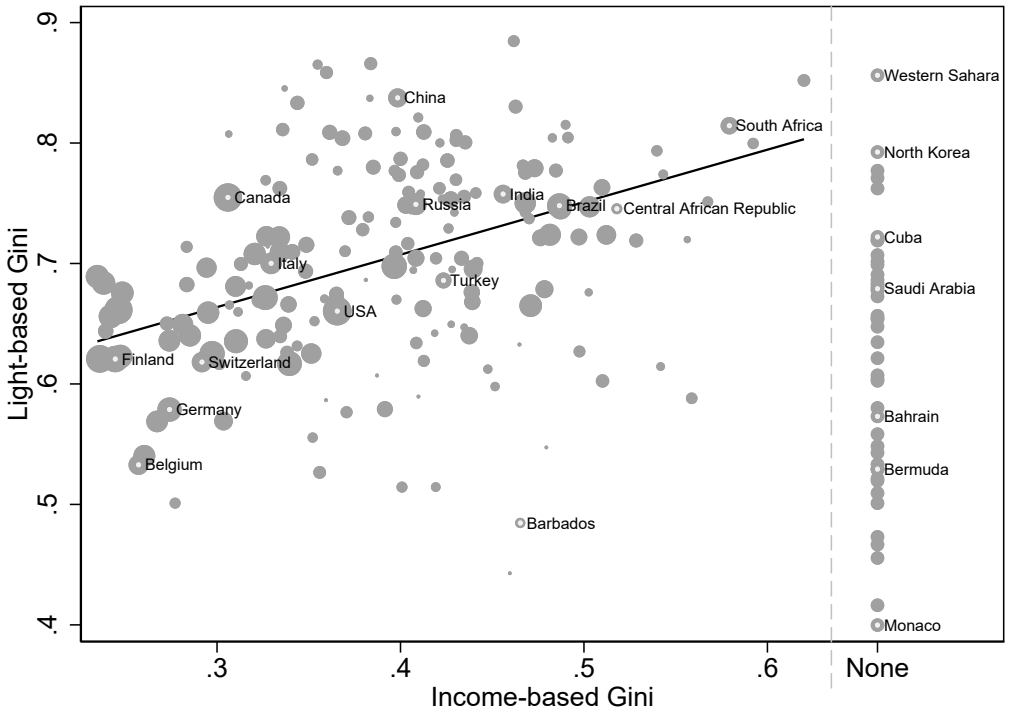

Notes: The graph shows country averages over all observed years. The size of the marker refers to the quality of the income-based Gini-coefficient, measured by the (sum of) inverse standard deviations across imputations. The depicted light-based Gini-coefficients refer to the weighted measures obtained with $\lambda = 0.5$. The right part of the graph depicts light-based averages for countries with no income-based Gini-coefficient data. Not shown: Norfolk Island (average light-based Gini-coefficient of 0.348) and Tokelau (average light-based Gini-coefficient of 0.235).

the (pooled) correlation between the two measures equals 0.9620 when using $\lambda = 0.5$. Given these results, we decided to rely on the SWIID-based Gini-coefficient due to the much larger number of observations. The following results will use the SWIID-based Gini-coefficient.

5 Applications

The aim of this section is to present three applications on different topics, namely out-of-pocket health care expenditure, epidemics and financial liberalization, using the geospatial

inequality measures discussed above. The main purpose is to show how the light- and income-based inequality measures correlate with different determinants of inequality and what can be learned from such correlations in terms of what types of economic activity are captured by both measures.

5.1 Empirical approach

In the following applications, we follow a unified approach, using the measures of inequality discussed above to estimate

$$G_{c,t} = \gamma z_{c,t} + \delta_t + \alpha_c + \epsilon_{c,t}, \tag{5}$$

where $G_{c,t}$ represents different versions of the Gini-coefficient (income- or light-based) for country c in year t. For comparative purposes, the Gini-coefficients are normalized by their corresponding sample mean and standard deviation. $z_{c,t}$ is the variable of interest that varies according to the application, and δ_t and α_c are time and country fixed effects, respectively. We also consider specifications without country fixed effects, in which cases we include country total areas as an explanatory variable to control for the negative correlation between the light-based inequality measures and country area, as documented in Figure 4 (a) – this has only minor quantitative effects on our estimates, see Table D4 of Appendix D. Finally, in order to account for the (potentially) unbalanced effect of less populated countries in our estimates, we present both unweighted and population-weighted regression results for all applications.

5.2 Out-of-pocket health care expenditure

In this application, we study the relationship between out-of-pocket (OOP) health expenditures and economic inequality. High OOP health care expenditures can result in financial hardship and can constitute a source of poverty in many developing countries (Xu et al., 2018). Alam and Mahal (2014) provide a literature review on this topic, concluding that OOP health spending poses a high burden on households' finances. Moreover, they find

that negative health shocks are often associated with significant reductions in labor supply, which can further aggravate the financial situation of households. In line with these results, a recent study by Christopher et al. (2018) finds that OOP expenditures increase inequality in the US.

We use data on the share of OOP health expenditures from the Global Health Expenditure Database (GHED) of the WHO (2019). These data are available from 2000 for 188 countries and territories, and contain the share of total health care expenditures coming out of households' pockets. The other two contributors observed in the data are health expenditures financed through the government budget and insurance contributions. While on average OOP expenditures make up 35% of total expenditures (see Table D3 in the Appendix), there is substantial variation across countries, ranging from a minimum of zero OOP health expenditures, observed for Belgium in the early 2000's, to a maximum of over 90% of health expenditures coming out-of-pocket in Myanmar in 2005.

Using these data, we estimate Equation (5), where $z_{c,t}$ is OOP expenditure. The results are provided in Table 2. OOP expenditure has a significant and positive relationship with the income-based Gini-coefficient (see column 1), which is stronger after weighting by population of the respective countries (column 4). Using the light-based inequality measure, we obtain a slightly larger coefficient estimate, which also increases when weighting by population. In terms of magnitude, our baseline results (column 1) indicate that a country with a 10 percentage point higher share of health expenditures coming from OOP has an expected income inequality about 0.07 standard deviations higher, while the same country's light-based inequality would be 0.10 standard deviations higher.

These results might be partly driven by the fact that individuals living in countries with more poverty and inequality also face a higher share of OOP expenditures. In order to account for this (potential) effect of reverse causality, we include country-specific fixed effects to the model in columns (2) and (5). As expected, including fixed effects leads to considerably smaller estimates, indicating that OOP expenditures have a less clear effect

Table 2: Regression Results for Out-of-Pocket Expenditure

Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)
Income-based Gini	0.735*	0.145		1.463***	-0.091	
	(0.403)	(0.120)		(0.541)	(0.212)	
Light-based Gini	1.002***	0.066	-0.004	1.779***	0.589**	0.537*
	(0.279)	(0.144)	(0.137)	(0.473)	(0.292)	(0.311)
Observations	2089	2089	2601	2089	2089	2601
# of countries	177	177	188	177	177	188
Country fixed effects	no	yes	yes	no	yes	yes
Population weights	no	no	no	yes	yes	yes

Notes: Each coefficient denotes one estimation of the model in Equation (5) using OOP expenditure as explanatory variable, and income- or light-based Gini-coefficients as dependent. Data on out-of-pocket (OOP) expenditure are from WHO (2019) and cover the years 2000-2013. The light-based Gini-coefficient refers to the weighted measure with our preferred parameter choice of $\lambda = 0.5$. Results for alternative values of λ are reported in Figure D1 (a) in the Appendix. The income-based Gini-coefficient is obtained from Solt (2016). Columns (1) and (4) include the country area as control variables. Columns (4)-(6) show results with population weights, where countries with higher population sizes receive a higher weight. *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered at the country level and reported in parentheses.

on the country-specific dynamics of inequality. Interestingly, the results using income-based inequality measures, both weighted and unweighted, turn insignificant, while the result for the light-based Gini-coefficient with population weights remains significant. Notice the latter finding is also corroborated using the extended sample of all available observations for the light-based measure, depicted in columns (3) and (6). These estimates can now be interpreted in terms of changes within a country, e.g.: using the estimates from column (5), a 30 p.p. decrease in OOP expenditure (approximately equivalent to the interquartile range of OOP shares) predicts a decrease of about 0.18 standard deviations in light-based inequality, while no statistically significant change can be identified solely in terms of income inequality. Finally, the estimated relationship with less OOP expenditure leading to lower (light-based) inequality fits the findings of the previous literature discussed above.

A potential explanation for these contrasting results is that traditional income inequality measures might be ill-suited for measuring inequality arising from health care spending (Kaestner and Lubotsky, 2016). Moreover, as already discussed above, the light-based

inequality measure is likely to capture more than income inequality. Government policies that increase universal health insurance coverage, and thereby reduce OOP expenditures, might not directly affect the distribution of incomes within a country, but rather affect the distribution and composition of households expenditures on light-producing goods and services. If OOP expenditures take up a larger share of income for the poorer households than for the richer, removing that burden can have a greater impact on the spending of the former, thereby affecting the distribution of expenditures. Besides, the overall allocation of expenditures between consumables and investments can also be affected if poorer and richer households allocate their spending differently. Considering that investments in physical structures, such as housing, are likely to increase the emission of lights more strongly, a policy reducing the OOP expenditures can have an effect in the light-based inequality measure that is not captured by the income-based measure. The fact that we find a positive and significant relationship between OOP and the light-based measure of inequality after controlling for country fixed effects suggests that reducing OOP expenditures can reduce economic inequality through such expenditure-composition channels.

5.3 Epidemics

The second application studies the relationship between epidemics and income- and light-based inequality measures. The economic literature analyzing this relationship is also limited. Cogneau and Grimm (2008) suggest that AIDS-induced mortality reduces total household income in Côte d'Ivoire, but find no distributional effects. Karlsson et al. (2014) on the other hand find an increase in poverty after the 1918 Spanish flu in Sweden. Another strand of literature studies how initial socio-economic inequality determines HIV outcomes, e.g.: Durevall and Lindskog (2012) find that young women living in neighborhoods with increasing wealth inequality face a higher risk of HIV infection.

For this application, we use data on epidemic disasters from the Emergency Events Database provided by the Centre for Research on the Epidemiology of Disasters (CRED,

2019). The database records a disaster event if at least one of the following criteria are fulfilled: 10 or more deaths, 100 or more people affected, a country declares a state of emergency or calls for international assistance. In the context of Equation (5), we use the percentage of individuals within each country that were affected by an epidemic during the year as the explanatory variable. Although the overall level and variation in this variable are relatively low, with both mean and standard deviation lower than 1% (see Table D3 in Appendix), its maximum value was registered by Kenya in 1994, when 26% of the country's population was affected by a parasitic disease epidemic.

Table 3 presents the estimation results. The estimated overall relationships between inequality and epidemics, reported in columns (1) and (4), are positive and statistically significant. The magnitude of these overall effects do not vary significantly across the inequality measures. Using the unweighted estimates, an epidemic affecting 10% of a country's population would be associated with an increase in inequality between 0.67 (light-based measure) and 0.69 (income-based measure) standard deviations. Including fixed effects leads to smaller, but still significant estimates for the income-based Gini-coefficient. However, the estimates for the light-based Gini-coefficient switch sign and become insignificant, suggesting that epidemics have no effect on the within country evolution of economic inequality.

These results can again be interpreted as evidence that the light-based measure captures multiple aspects of economic inequality that go beyond income effects. Particularly in the context of epidemics, the associated income shocks are likely to be of a transitory nature if the epidemic has little effect on productive means, while the light-based measures can capture more persistent effects that a regular exposition to epidemics can have on infrastructure and development across countries. Finally, extending the sample to all observations jointly available in the CRED (2019) database, as reported in columns (3) and (6), leads to more than 5,000 observations across the 234 countries in our sample. In spite of the stark contrast in sample sizes, evidencing the data limitations of income-based Gini-coefficients, the within country effect of epidemics on the light-based inequality measure remains insignificant.

Table 3: Regression Results for Epidemics

Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)
Income-based Gini	6.903**	0.332**		6.246***	0.829*	
	(2.750)	(0.147)		(2.201)	(0.446)	
Light-based Gini	6.710***	-0.328	-1.065	5.749**	-0.264	-0.245
	(1.562)	(0.721)	(0.817)	(1.185)	(0.367)	(0.432)
Observations	3278	3278	5148	3278	3278	5148
# of countries	187	187	234	187	187	234
Country fixed effects	no	yes	yes	no	yes	yes
Population weights	no	no	no	yes	yes	yes

Notes: Each coefficient denotes one estimation of the model in Equation (5) using the percentage of individuals affected by an epidemic as explanatory variable, and income- or light-based Gini-coefficients as dependent. Data on epidemics are from CRED (2019) and cover the years 1992-2013. The light-based Gini-coefficient refers to the weighted measure with our preferred parameter choice of $\lambda = 0.5$. Results for alternative values of λ are reported in Figure D1 (b) in the Appendix. The income-based Gini-coefficient is obtained from Solt (2016). Columns (1) and (4) include the country area as control variables. Columns (4)-(6) show results with population weights, where countries with higher population sizes receive a higher weight. *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered at the country level and reported in parentheses.

5.4 Capital account liberalization

In the last application, we study the relationship between capital account liberalization and inequality. Bumann and Lensink (2016) show that liberalization can lead to lower income inequality for countries with a sufficiently high level of financial depth. De Haan and Sturm (2017) review the empirical literature and report new evidence that financial liberalization tends to increase income inequality, though this result is also conditioned on the quality of political institutions and the level of financial development. Similarly, Furceri and Loungani (2018) find that episodes of capital account liberalization increase income inequality, particularly in countries that lack financial depth.

Following this literature, we use the Chinn and Ito (2008) index of capital account openness as our explanatory variable of interest. The index is constructed through an analysis of principal components of data on restrictions to cross-border financial transactions, the latter collected from the IMF's Annual Report on Exchange Arrangements and Exchange

Restrictions. Although the principal components approach renders a measurement without a direct economic interpretation, the index provides a useful criterion for relative assessments of capital account liberalization across countries and over time. It reflects, for example, the relatively higher levels of financial openness of industrialized countries, such as the United Kingdom with an average index value of 1.35 standard deviations above the sample mean; it also captures well known cases of rapid financial liberalization, as in the case of Brazil jumping from the 6th to the 58th percentile of the index sample between 1997 and 2006.

Table 4 provides the estimation results. Without fixed effects, in columns (1) and (4), the relationship between capital account liberalization and inequality is significantly negative, indicating that liberalized countries tend to be more equal. These results, however, can be due to effects on both directions, i.e., that countries with lower inequality tend to liberalize more at the same time that greater financial openness can decrease inequality. In an attempt to clarify this effect, we include country fixed effects, which allow us to filter out the cross-country relationship and focus on the within country dynamics. Once we include country fixed effects, financial openness is found to have a positive and significant effect on the light-based Gini-coefficient, a result that is robust to both the sample definition, and the estimation with population weights. Hence, these results indicate that countries which experienced episodes of capital account liberalization also tended to endure periods of increased inequality, which is in line with earlier literature. In terms of effect size the results in columns (5) and (6) indicate that an increase in financial openness from the 25th to the 75th percentile (equivalent to an index increase of 3.3; see Table D3 in the Appendix) is associated with a 0.12 standard deviations increase in inequality.

It is interesting to note that this effect is only captured by the light-based inequality measure, while the income-based measure is either insignificantly related with financial liberalization episodes, under unweighted estimation, or presenting a negative and (marginally) significant relationship when weighting by population. This clearly illustrates another case where the focus on income data can have different, and potentially misleading,

Table 4: Regression Results for Capital Account Liberalization

Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)
Income-based Gini	-0.186***	0.011		-0.206***	-0.079*	
	(0.043)	(0.013)		(0.059)	(0.043)	
Light-based Gini	-0.152***	0.040***	0.031***	-0.272***	0.036*	0.034*
	(0.032)	(0.011)	(0.011)	(0.063)	(0.018)	(0.017)
Observations	3013	3013	3780	3013	3013	3780
# of countries	170	170	181	170	170	181
Country fixed effects	no	yes	yes	no	yes	yes
Population weights	no	no	no	yes	yes	yes

Notes: Each coefficient denotes one estimation of the model in Equation (5) using the Chinn-Ito index as explanatory variable, and income- or light-based Gini-coefficients as dependent. Data on financial openness are from Chinn and Ito (2008) and cover the years 1992-2013. The light-based Gini-coefficient refers to the weighted measure with our preferred parameter choice of $\lambda = 0.5$. Results for alternative values of λ are reported in Figure D1 (c) in the Appendix. The income-based Gini-coefficient is obtained from Solt (2016). Columns (1) and (4) include the country area as control variables. Columns (4)-(6) show results with population weights, where countries with higher population sizes receive a higher weight. *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered at the country level and reported in parentheses.

implications. In particular, the estimates in column (5) indicate that while financial liberalization policies may help reduce income inequality in more populated countries, such policies have a clearer and more robust effect of increasing inequality in a broader sense, as captured by the light-based measure. Once financial markets open up, tax evasion and shadow economy activities are likely to increase. Traditional income inequality measures, relying on declared income, are unable to take such activities into account (Alstadsæter et al., 2019), while the light-based measure seems to capture them. We interpret these results as further evidence that the light-based Gini-coefficients can provide a more robust measurement for the identification of policy effects on economic inequality.

5.5 Results summary

In the applications above, we have studied how inequality is related to health expenditures, epidemics, and financial liberalization, across global samples of countries. We focused on how

the light-based measure of inequality developed in this paper compares to the traditional income-based Gini-coefficient. Interestingly, the two measures of inequality generally agreed on the overall relationship estimates: inequality tends to be higher in countries with higher out-of-pocket (OOP) health expenditures, as well as in countries with a higher incidence of epidemics, while financial liberalization is higher in more equal countries. Nevertheless, these relationships are mainly driven by cross-country variation.

In fact, the two measures lead to contrasting results with respect to country-specific dynamics (i.e., after including country fixed-effects in the regressions). First, we find a positive and significant effect of OOP health expenditures on light-based inequality, but not on the income-based measure of inequality. Second, while the effect of epidemics on income inequality remains positive, there is no significant effect on light-based inequality. Finally, financial openness is found to have a positive and robust within country effect on light-based inequality, but not on income inequality.

Our interpretation of these differences goes back to the rationale for using night lights data as a broader measurement of economic prosperity than income. Because the lights data can capture other dimensions of economic activity, such as informal activities, infrastructure investments, consumption, and wealth, it is natural that we obtain different implications for the within country dynamics of inequality. Besides, one can not rule out the fact that the geographical nature of our source data imprint stronger traces of regional inequality in the geospatial measures. Altogether, we interpret these results as evidence that our new measure of inequality is less prone to the effects of transitory income shocks, hence the insignificant effect of epidemics, at the same time that it can capture richer dynamics in the more structural contexts of health expenditures and financial development.

Finally, we also checked for the robustness of these estimates to our choice of $\lambda = 0.5$ for the construction of the weighted inequality measure. Considering how this parameter affected the cross- and within-country correlations between the light-based and the income-based measures, its calibration could affect the comparative between the results obtained

with these measures. Estimation results for alternative values of λ are presented in Figure D1 in the Appendix. Overall, we find little differences in estimates for $\lambda > 0.1$. As expected, as $\lambda \to 0$, the commonality between the pooled estimates obtained with the two measures is lost; on the other hand, the estimates including country fixed-effects are hardly affected, except for a narrowing of confidence bands.

6 Concluding Remarks

In this paper, we construct a measure of economic inequality based on geospatial nighttime lights and population data. The underlying assumption is that nighttime light emissions are a proxy for economic prosperity. By looking at how lights and population are distributed geographically, we gauge multiple aspects of economic inequality. In particular, we argue that our new measure can capture broad features of economic activity within countries, such as the distribution of income, consumption expenditures, infrastructure investments, and wealth. Besides, in comparison to traditional sources of economic data, night lights can more promptly capture the effects of informal economic activities.

The use of night lights data for the measurement of inequality is not entirely new to the literature. After reviewing this literature, we discuss our methodological novelty. Some key innovations in our paper include an account for alternative sources of data and levels of aggregation, as well as the development of a calibration approach geared towards the measurement of income inequality. Besides these methodological advances, we also consider a broader sample than the previous literature, covering a total of 234 countries and territories between 1992 and 2013. The new measure is significantly correlated with cross-country variation in income inequality. We also find that our light-based measures suggest generally higher inequality than traditional measures.

In addition, we provide three applications of the data, which allow us to compare the results between income- and light-based measures. These applications cover different

fields of economics, such as health economics and international finance. In all applications, we find similar results for both measures in the cross-country estimations but important differences in terms of country-specific dynamics of inequality. We interpret these differences as evidence that our new measure of inequality is less prone to the effects of transitory income shocks, while capturing richer dynamics of more deeply rooted structural factors, such as the distribution of productive means and the composition of expenditures across economic classes.

Naturally, the geospatial inequality measures have also some caveats, as both the night lights and population data have their respective measurement issues. For instance, the quality of measurement by satellites can change considerably over time, while there is substantial cross-country variation in the accuracy of the population data. Another important issue relates to the sensitivity of the remote sensing instruments to capture lights from dimmer locations that can be inhabited by the most deprived population. In the paper, we strive to account for these drawbacks and their potential to affect our measures by following the latest literature and designing new solutions. In particular, our approach is geared towards an agnostic use of all information that is readily available from these rich sources of data, enabling us to obtain more accurate measures of inequality.

Further research can employ this novel data to study important policy questions related to inequality. This might be especially relevant in developing countries where no or only a few datapoints are available when using traditional inequality data. Moreover, our methodology can be easily extended to construct inequality measures at a subnational level. Finally, one could go beyond the Gini-coefficient and construct other inequality measures relying on nighttime lights data.

Data Access

The data on light-based inequality measures generated by this project are available at: `http://www.ciesin.columbia.edu/data/global-geospatial-inequality`.

References

Acemoglu, D., S. Naidu, P. Restrepo, and J. A. Robinson (2015). Chapter 21 – Democracy, Redistribution, and Inequality. In A. B. Atkinson and F. Bourguignon (Eds.), *Handbook of Income Distribution*, Volume 2, pp. 1885 – 1966. Elsevier.

Alam, K. and A. Mahal (2014). Economic impacts of health shocks on households in low and middle income countries: a review of the literature. *Globalization and Health 10*, 21.

Alesina, A., S. Michalopoulos, and E. Papaioannou (2016). Ethnic Inequality. *Journal of Political Economy 124*(2), 428–488.

Alstadsæter, A., N. Johannesen, and G. Zucman (2019). Tax evasion and inequality. *American Economic Review 109*(6), 2073–2103.

Atkinson, A. B. and A. Brandolini (2001). Promise and Pitfalls in the Use of "Secondary" Data-Sets: Income Inequality in OECD Countries As a Case Study. *Journal of Economic Literature 39*(3), 771–799.

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

Bickenbach, F., E. Bode, P. Nunnenkamp, and M. Söder (2016). Night lights and regional GDP. *Review of World Economics 152*, 425–447.

Bluhm, R. and M. Krause (2018). Top Lights – Bright Cities and Their Contribution to Economic Development. *CESifo Working Paper No. 7411*.

Bright, E. A., A. N. Rose, M. L. Urban, and J. L. McKee (2017). *LandScan 2016*. Oak Ridge, TN: Oak Ridge National Laboratory.

Browning, M. and T. Crossley (2009). Are two cheap, noisy measures better than one expensive, accurate one? *The American Economic Review 99*(2), 99–103.

Bumann, S. and R. Lensink (2016). Capital account liberalization and income inequality. *Journal of International Money and Finance 61*, 143–162.

Chen, X. and W. D. Nordhaus (2011). Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences 108*(21), 8589–8594.

Chinn, M. D. and H. Ito (2006). What matters for financial development? Capital controls, institutions, and interactions. *Journal of Development Economics 81*(1), 163–192.

Chinn, M. D. and H. Ito (2008). A New Measure of Financial Openness. *Journal of Comparative Policy Analysis: Research and Practice 10*(3), 309–322.

Christopher, A. S., D. U. Himmelstein, S. Woolhandler, and D. McCormick (2018). The Effects of Household Medical Expenditures on Income Inequality in the United States. *American Journal of Public Health 108*(3), 351–354.

CIESIN (2016). Documentation for Gridded Population of the World, Version 4 (GPWv4). Technical report, Center For International Earth Science Information Network – Columbia University, Palisades, NY: NASA Socioeconomic Data and Applications Center (SEDAC).

Cogneau, D. and M. Grimm (2008). The Impact of AIDS Mortality on the Distribution of Income in Côte d'Ivoire. *Journal of African Economies 17*(5), 688–728.

CRED (2019). Emergency Events Database (EM-DAT). Centre for Research on the Epidemiology of Disasters. https://www.emdat.be/. Accessed 06.03.2020.

Davies, J., S. Sandström, A. Shorrocks, and E. Wolff (2009). The global pattern of household wealth. *Journal of International Development 21*, 1111–1124.

De Haan, J. and J. E. Sturm (2017). Finance and income inequality: A review and new evidence. *European Journal of Political Economy 50*, 171–195.

Deaton, A. (2003). Health, Inequality, and Economic Development. *Journal of Economic Literature 41*(1), 113–158.

Deaton, A. (2005). Measuring Poverty in a Growing World (or Measuring Growth in a Poor World). *The Review of Economics and Statistics 87*(1), 1–19.

Deaton, A. (2016). Measuring and Understanding Behavior, Welfare, and Poverty. *American Economic Review 106*(6), 1221–1243.

Deltas, G. (2003). The Small-Sample Bias of the Gini Coefficient: Results and Implications for Empirical Research. *The Review of Economics and Statistics 85*(1), 226–234.

Dobson, J. E., E. A. Bright, P. R. Coleman, R. C. Durfee, and B. A. Worley (2000). LandScan: A Global Population Database for Estimating Populations at Risk. *Photogrammetric Engineering & Remote Sensing 66*(7), 849–857.

Donaldson, D. and A. Storeygard (2016). The view from above: Applications of satellite data in economics. *Journal of Economic Perspectives 30*(4), 171–98.

Doxsey-Whitfield, E., K. MacManus, S. B. Adamo, L. Pistolesi, J. Squires, O. Borkovska, and S. R. Baptista (2015). Taking Advantage of the Improved Availability of Census Data: A First Look at the Gridded Population of the World, Version 4. *Papers in Applied Geography 1*(3), 226–234.

Durevall, D. and A. Lindskog (2012). Economic Inequality and HIV in Malawi. *World Development 40*(7), 1435–1451.

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

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

Elvidge, C. D., V. R. Hobson, I. L. Nelson, J. M. Safran, B. T. Tuttle, J. B. Dietz, and K. E. Baugh (2003). Overview of DMSP OLS and scope of applications. In V. Mesev (Ed.), *Remotely Sensed Cities*, pp. 281–300. Taylor & Francis.

Elvidge, C. D., J. Safran, I. L. Nelson, B. T. Tuttle, V. R. Hobson, K. E. Baugh, J. B. Dietz, and E. H. Erwin (2004). Area and positional accuracy of DMSP nighttime lights data. In *Remote Sensing and GIS Accuracy Assessment*, pp. 281–292. CRC Press.

Elvidge, C. D., D. Ziskin, K. E. Baugh, B. T. Tuttle, T. Ghosh, D. W. Pack, E. H. Erwin, and M. Zhizhin (2009). A Fifteen Year Record of Global Natural Gas Flaring Derived from Satellite Data. *Energies 2*(3), 595–622.

Ferreira, F. H. G., N. Lustig, and D. Teles (2015). Appraising cross-national income inequality databases: An introduction. *Journal of Economic Inequality 13*, 497–526.

Furceri, D. and P. Loungani (2018). The distributional effects of capital account liberalization. *Journal of Development Economics 130*, 127–144.

Galbraith, J. K. (2019). Sparse, Inconsistent and Unreliable: Tax Records and the World Inequality Report 2018. *Development and Change 50*(2), 329–346.

Galimberti, J. K. (2020). Forecasting GDP growth from outer space. *Oxford Bulletin of Economics and Statistics*, in press.

Gibson, J., S. Olivia, and G. Boe-Gibson (2020). Night lights in economics: Sources and uses. *Journal of Economic Surveys in press*.

Henderson, J. V., T. Squires, A. Storeygard, and D. Weil (2018). The Global Distribution of Economic Activity: Nature, History, and the Role of Trade. *The Quarterly Journal of Economics 133*(1), 357–406.

Henderson, J. V., A. Storeygard, and D. N. Weil (2012). Measuring Economic Growth from Outer Space. *American Economic Review 102*(2), 994–1028.

Hodler, R. and P. A. Raschky (2014). Regional Favoritism. *Quarterly Journal of Economics 129*(2), 995–1033.

Hsu, F.-C., K. E. Baugh, T. Ghosh, M. Zhizhin, and C. D. Elvidge (2015). DMSP-OLS radiance calibrated nighttime lights time series with intercalibration. *Remote Sensing 7*(2), 1855–1876.

Hu, Y. and J. Yao (2019). Illuminating Economic Growth. *IMF Working Papers 19/77*.

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

Kaestner, R. and D. Lubotsky (2016, May). Health insurance and income inequality. *Journal of Economic Perspectives 30*(2), 53–78.

Karlsson, M., T. Nilsson, and S. Pichler (2014). The impact of the 1918 Spanish flu epidemic on economic performance in Sweden: An investigation into the consequences of an extraordinary mortality shock. *Journal of Health Economics 36*, 1–19.

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

Mellander, C., J. Lobo, K. Stolarick, and Z. Matheson (2015). Night-Time Light Data: A Good Proxy Measure for Economic Activity? *PLoS ONE 10*(10), 1–18.

Michalopoulos, S. and E. Papaioannou (2013). Pre-Colonial Ethnic Institutions and Contemporary African Development. *Econometrica 81*(1), 113–152.

Milanovic, B. (2005). Can We Discern the Effect of Globalization on Income Distribution? Evidence from Household Surveys. *World Bank Economic Review 19*(1), 21–44.

Miller, S. D., W. Straka, S. P. Mills, C. D. Elvidge, T. F. Lee, J. Solbrig, A. Walther, A. K. Heidinger, and S. C. Weiss (2013). Illuminating the Capabilities of the Suomi National Polar-Orbiting Partnership (NPP) Visible Infrared Imaging Radiometer Suite (VIIRS) Day/Night Band. *Remote Sensing 5*(12), 6717–6766.

Mveyange, A. (2015). Night lights and regional income inequality in Africa. *WIDER Working Paper 2015/085*.

NOAA (2017). Image and Data processing by NOAA's National Geophysical Data Center. DMSP data collected by the US Air Force Weather Agency. Version 4. https://ngdc.noaa.gov/eog/download.html, last accessed May 29, 2020.

Nordhaus, W. and X. Chen (2015). A sharper image? Estimates of the precision of nighttime lights as a proxy for economic statistics. *Journal of Economic Geography 15*(1), 217–246.

Piketty, T. and E. Saez (2014). Inequality in the long run. *Science 344*(6186), 838–843.

Pinkovskiy, M. and X. Sala-i Martin (2016). Lights, Camera . . . Income! Illuminating the National Accounts-Household Surveys Debate. *Quarterly Journal of Economics 131*(2), 579–631.

Sen, A. K. (1997). From Income Inequality to Economic Inequality. *Southern Economic Journal 64*(2), 383–401.

Shorrocks, A. and G. Wan (2005). Spatial decomposition of inequality. *Journal of Economic Geography 5*(1), 59–81.

Solt, F. (2016). The Standardized World Income Inequality Database. *Social Science Quarterly 97*(5), 1267–1281. SWIID Version 7.1, August 2018.

UNU-WIDER (2017). World Income Inequality Database (WIID 3.4). https://www.wider.unu.edu/project/wiid-world-income-inequality-database, last accessed May 29, 2020.

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

WHO (2019). Global Health Expenditure Database. http://apps.who.int/nha/database, last accessed March 6, 2020.

Xu, K., A. Soucat, J. Kutzin, C. Brindley, E. Dale, N. Van der Maele, T. Roubal, C. Indikadahena, H. Toure, and V. Cherilova (2018). New Perspectives on Global Health Spending for Universal Health Coverage. Technical report, World Health Organization.

Appendix for Online Publication

A Population Sample Extrapolation

In order to construct inequality measures for the broader sample of data available on night lights, we extrapolate the population data to cover an annual sample from 1992 to 2013. For that purpose we calculate GPW population growth rates at the pixel level (which are effectively varying only at the census level) and use these growth rates to extrapolate data for the different sources. In fact, we conduct these calculations retroactively, starting with the interpolation of the GPW data over the more recent missing middle years, and then computing backcasts of both GPW and LSC data for the earlier years. In particular, letting $p^{\mathcal{S}}_{i,t}$ stand for the population count observed in pixel i, for year t, and using source data $\mathcal{S} = \{GPW, LSC\}$, we first calculate

$$\Delta_{i,2015/2010} = \frac{p^{GPW}_{i,2015}}{p^{GPW}_{i,2010}} \,, \tag{A1}$$

for every pixel in the global population dataset. We then obtain GPW interpolated annual population counts for the years 2011 to 2013 using

$$\hat{p}^{GPW}_{i,2010+n} = p^{GPW}_{i,2010} \times \Delta^{\frac{n}{5}}_{i,2015/2010} \,, \tag{A2}$$

for $n = \{1, 2, 3\}$. We repeat this procedure for the years of 2006-2009, 2001-2004, adjusting the calculation of (A1) and (A2) according to the subperiod. We then use the earliest change ratio, $\Delta_{i,2005/2000}$, to calculate backcasts for both GPW and LSC according to

$$\hat{p}^{\mathcal{S}}_{i,2000-n} = p^{\mathcal{S}}_{i,2000} \times \Delta^{-\frac{n}{5}}_{i,2005/2000} \,, \tag{A3}$$

for $n = \{1, \ldots, 8\}$. Notice our use of GPW growth rates as well for backcasting LSC population counts is due to the greater reliability of the former for comparisons across time. We have indeed experimented with LSC growth rates but their implied population estimates tended to be highly unstable.

B Gini-coefficient Derivation

The Gini-coefficient is defined as the ratio of the area that lies between the line of perfect equality (45° line) and the Lorenz curve over the total area under the equality line. For a discrete population, obtaining the first area is equivalent to summing up the difference between the equality line and the individual lights as a fraction of the total. Namely, individual i's area is given by

$$\frac{1}{N}\left(\frac{i}{N}-\frac{\sum_{k=1}^{i} y_k}{\sum_{k=1}^{N} y_k}\right),$$

where N is the total number of people and y_i is the light of individual i (notice a slight abuse of notation as i here should not be confused with its use in the main text to denote a pixel in the gridded dataset); geometrically, the expression inner brackets corresponds to the height of the individual's income "bar", while $1/N$ is its width. Summing up over all individuals we obtain

$$\begin{aligned}
A &= \frac{1}{N}\sum_{i=1}^{N}\left(\frac{i}{N}-\frac{\sum_{k=1}^{i} y_k}{\sum_{k=1}^{N} y_k}\right), \\
&= \frac{1}{N}\left(\frac{1}{N}\sum_{i=1}^{N} i-\sum_{i=1}^{N}\sum_{k=1}^{i} y_k/\sum_{k=1}^{N} y_k\right), \\
&= \frac{1}{N}\left(\frac{N+1}{2}-\sum_{i=1}^{N}(N-i+1)\,y_i/\sum_{i=1}^{N} y_i\right).
\end{aligned}$$

The Gini-coefficient is then obtained from its definition as the area A divided by 1/2 (area underneath the 45° line), hence

$$G' = \frac{1}{N}\left(N+1-2\sum_{i=1}^{N}(N-i+1)\,y_i/\sum_{i=1}^{N} y_i\right).$$

This formula can be further simplified as follows

$$\begin{aligned}
G' &= \frac{1}{N}\left(N+1-2\left((N+1)\sum_{i=1}^{N} y_i-\sum_{i=1}^{N} i y_i\right)/\sum_{i=1}^{N} y_i\right), \\
&= \frac{1}{N}\left(N+1-2\left(N+1-\sum_{i=1}^{N} i y_i/\sum_{i=1}^{N} y_i\right)\right), \\
&= \frac{2}{N}\frac{\sum_{i=1}^{N} i y_i}{\sum_{i=1}^{N} y_i}-\frac{N+1}{N}. \qquad \text{(B4)}
\end{aligned}$$

Grouped data In order to obtain the Gini-coefficient for grouped data on population and light intensity we need to adjust Equation (B4) to account for variable population sizes across

the grouped data points. Let M denote the total number of groups, n_j be the population in group j, and x_j the light intensity in group j, with the index ranked in ascending order of groups' light intensity per capita. Expanding the summation in the numerator of (B4) we find that

$$\begin{aligned}\sum_{i=1}^{N} iy_i &= (1+\ldots+n_1)\frac{x_1}{n_1} + (n_1+1+\ldots+n_2)\frac{x_2}{n_2} + \ldots \\ &\quad \ldots + (n_{M-1}+1+\ldots+n_M)\frac{x_M}{n_M}, \\ &= \sum_{j=1}^{M} \mathcal{S}(a_{j-1}+1, n_j)\frac{x_j}{n_j}, \end{aligned} \tag{B5}$$

where $\mathcal{S}(p,q)$ stands for the sum of the sequence of integers from p to $p+q$, which is given by $\mathcal{S}(p,q) = \frac{q}{2}(2p+q-1)$, and a_j is the cumulative sum of population up to group j, i.e., $a_j = \sum_{k=1}^{j} n_k$ for $j > 0$, and $a_0 = 0$. Substituting these quantities in (B5) we obtain

$$\begin{aligned}\sum_{i=1}^{N} iy_i &= \sum_{j=1}^{M} \frac{n_j}{2}\left(2(a_{j-1}+1)+n_j-1\right)\frac{x_j}{n_j}, \\ &= \sum_{j=1}^{M}\left(\frac{n_j+1}{2}+a_{j-1}\right)x_j,\end{aligned}$$

which can then be substituted back into (B4) to obtain the grouped data formulation of the Gini-coefficient,

$$G = \frac{2}{N}\frac{\sum_{j=1}^{M}\left(\frac{n_j+1}{2}+a_{j-1}\right)x_j}{\sum_{j=1}^{M} x_j} - \frac{N+1}{N}. \tag{B6}$$

Small-sample bias Calculated from discrete data the Gini-coefficient is downward biased due to the effects of interpolation on the coefficient's upper bound, particularly lowering this upper bound below unity. To see this consider first the case with individual data, Equation (B4), when all individuals except one have zero lights, i.e., $y_i = 0 \forall i < N$, then

$$\begin{aligned}G'_{max} &= \frac{2}{N}\frac{Ny_N}{y_N} - \frac{N+1}{N}, \\ &= \frac{N-1}{N} = 1 - \frac{1}{N}.\end{aligned}$$

Hence, the smaller the sample of data used for the calculation of the Gini-coefficient the lower its upper bound. A straightforward correction for this distortion is to divide G' by $1 - 1/N$ (Deltas, 2003). A similar result and correction is possible for the grouped data formulation, Equation (B6). Here, the hypothetical maximum Gini-coefficient is given by a situation where

$x_j = 0 \forall j < M$, which then results in

$$\begin{aligned} G_{max} &= \frac{2}{N} \frac{\left(\frac{n_M+1}{2} + a_{M-1}\right) x_M}{x_M} - \frac{N+1}{N}, \\ &= \frac{2}{N} \left(\frac{n_M+1}{2} + N - n_M \right) - \frac{N+1}{N}, \\ &= \frac{N - n_M}{N} = 1 - \frac{n_M}{N}. \end{aligned} \tag{B7}$$

In contrast to the calculation based on individual data, the upper bound on the Gini-coefficient calculated from grouped data decreases with the share of population in the group concentrating total lights. To translate this result into a correction it is important to recall that population shares can vary across groups, in which case the maximal inequality is in fact obtained in the hypothetical situation of all lights being concentrated in the group with minimal population share. Hence, a correction for the small-sample bias under grouped data is obtained by dividing G by $1 - \min\{n_j/N\}$. A simpler and stronger correction is obtained by assuming n_M is equal to the average of n_j, i.e., replacing $n_M = \frac{1}{M}\sum_{j=1}^{M} n_j = N/M$ in (B7), the correction simplifies to $1 - 1/M$. Notice this is equivalent to assuming equal population shares across groups.

C Summary Statistics on Gini-coefficients

Table C1 shows summary statistics on our geospatial Gini-coefficients. We can draw several observations from these statistics. First, we observe that the census-level measures of inequality are on average lower than the pixel-level ones. Second, the scaling factors are found to have a nonlinear effect on our inequality measures, initially lowering the average Gini-coefficients as κ increases from 0.5 to 2, and then increasing them as κ is calibrated above 2. Third, a similar nonlinearity affects the dispersion of our geospatial Gini-coefficients. Interestingly, the census-level measures are more dispersed than the pixel-level ones. One potential explanation for this difference is that the level of geographical aggregation in the GPW-based measures varies substantially across countries, according to the different census areas, whereas the pixel-level aggregation is nearly constant. Due to the Earth's curvature, the area covered by each pixel depends on its latitude, decreasing from the equator to the poles. Hence, countries located closer to the equator have datapoints aggregated over bigger areas than Northern and Southern countries. Finally, a decomposition of the variance of the geospatial Gini-coefficients is also reported in Table C1, showing that most of the dispersion in these measures is due to cross-country variation. Also notice that, in spite of having a lower overall variance, the LSC-based Gini-coefficients show greater variation over time than those based on GPW (except for $\kappa = 5$); this result corroborates our assessment on the time-varying precision of the LSC source.

Table C1: Geospatial Gini-Coefficients Summary Statistics

(a) LSC/Pixel-level Gini-Coefficients

	$\kappa = 0.5$	$\kappa = 1.0$	$\kappa = 1.5$	$\kappa = 2.0$	$\kappa = 3.0$	$\kappa = 5.0$
Averages	0.788	0.757	0.729	0.714	0.751	0.818
Standard deviations:						
Pooled observations	0.103	0.095	0.086	0.081	0.092	0.114
Across country averages	0.086	0.079	0.071	0.067	0.080	0.104
Across year averages	0.025	0.022	0.018	0.015	0.012	0.007

(b) GPW/Census-level Gini-Coefficients

	$\kappa = 0.5$	$\kappa = 1.0$	$\kappa = 1.5$	$\kappa = 2.0$	$\kappa = 3.0$	$\kappa = 5.0$
Averages	0.600	0.538	0.475	0.443	0.510	0.627
Standard deviations:						
Pooled observations	0.195	0.201	0.197	0.180	0.177	0.212
Across country averages	0.194	0.200	0.195	0.177	0.172	0.207
Across year averages	0.003	0.006	0.009	0.008	0.012	0.011

Notes: The pooled standard deviations are calculated over all 5,148 country-year geospatial Gini-coefficients together, whereas the standard deviations on country and year averages are calculated from 234 country and 22 year averages, respectively.

Table C2: Income- and Light-based Gini-coefficients Summary Statistics

		Weighted light-based Gini-coefficients					
	Income	Common sample			Extended sample		
	Ginis	$\lambda = 0$	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 0$	$\lambda = 0.5$	$\lambda = 1$
Averages	0.389	0.642	0.682	0.703	0.677	0.705	0.721
Standard deviations:							
Pooled observations	0.084	0.136	0.104	0.098	0.109	0.082	0.078
Across country averages	0.081	0.136	0.100	0.093	0.115	0.083	0.077
Across year averages	0.007	0.003	0.007	0.006	0.005	0.007	0.007

Notes: Statistics under the common sample for income- and weighted light-based Gini-coefficients are based on 3,278 country-year observations, while the extended sample contains a total of 5,148 observations.

Table C2 presents summary statistics on the light- and income-based Gini-coefficients. There is a clear difference in levels between the light- and the income-based Gini-coefficients, the former being between 0.28 and 0.33 higher than the latter, on average and depending on λ. The overall dispersion of inequality is also greater under the light-based measurements, though the light-based Gini-coefficients obtained from a higher weight on cross-country correlations ($\lambda \to 1$) tend to have similar standard deviation statistics as the income-based Gini-coefficients. This is consistent with our analysis of the geospatial Gini-coefficients variances above–recall most part of the light-based Gini-coefficients variation comes from their cross-country differences, which is also the case for the income- and light-based Gini-coefficients. It is also important to note that the light-based measures are highly correlated across their different λ's – their pooled correlations vary from 0.790 (between the two extreme λ values) to 0.999 (between smaller λ steps).

D Supplementary Statistics

Figure D1: Robustness Graphs for Applications

(a) Out-of-Pocket Health Care Expenditure

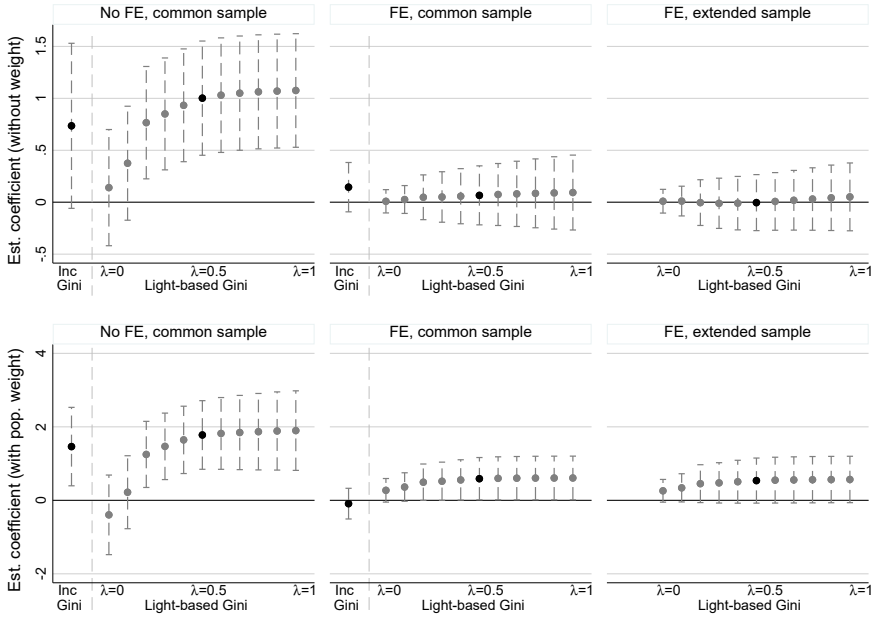

(b) Epidemics

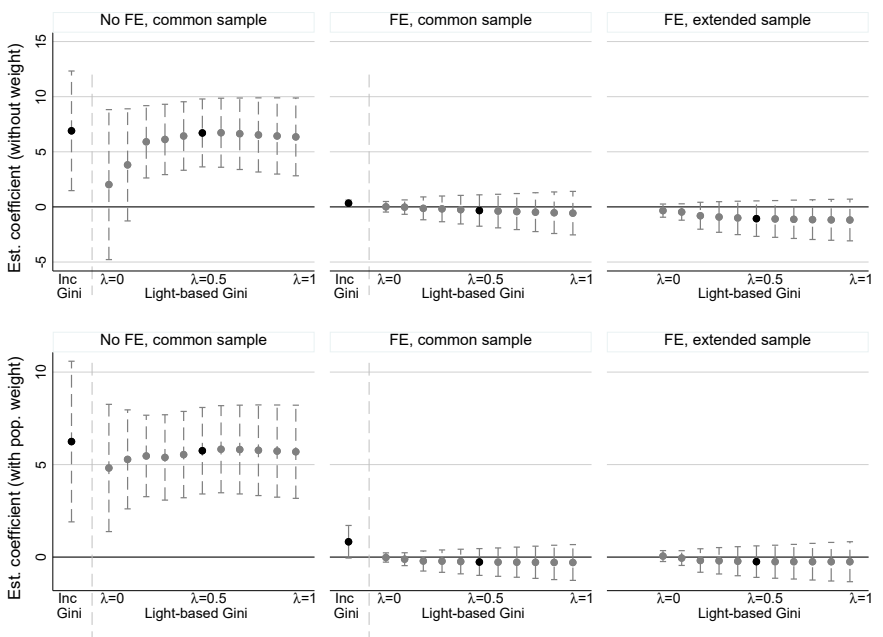

Figure D1: Robustness Graphs for Applications (Continued)

(c) Capital Account Liberalization

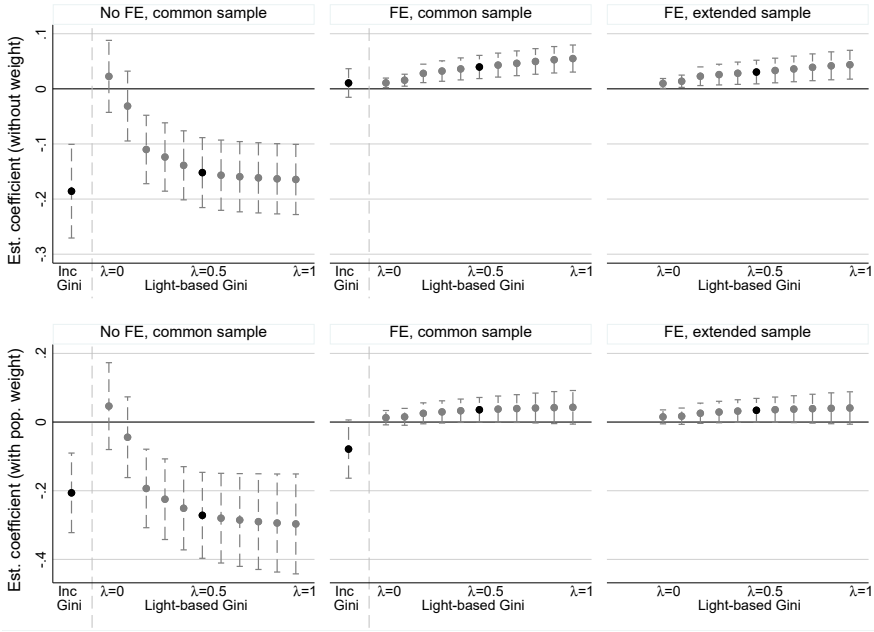

Notes: The graphs show point estimates and confidence intervals for estimating Equation (5). *Inc Gini* refers to the income-based Gini-coefficient as the dependent variable, while *Light-based Gini* stands for the light-based Gini-coefficient. We show results for different values of the λ parameter increasing in steps of 0.1 from 0 to 1. Estimates presented in Tables 2, 3, and 4 are in black. In particular, using the income-based Gini-coefficient as the dependent variable is represented by the first estimate and the estimate for $\lambda = 0.5$ may be found in the middle. Other values of λ are in gray. The explanatory variable is the share of OOP expenditure (Source: WHO, 2019) for (a), the percentage of individuals affected by an epidemic disaster (Source: CRED, 2019) for (b), and capital account liberalization (Chinn and Ito, 2006) for (c).

Table D3: Descriptive Statistics for Applications

Variable	**Obs**	**Mean**	**Std. Dev.**	**Min**	**P25**	**P50**	**P75**	**Max**
Share of OOP health exp	2601	.3503	.1999	0	.1813	.3304	.4992	.9081
Perc. of inh. affected by epi.	5148	.0004	.0061	0	0	0	0	.261
Capital account liberalization	3780	.232	1.56	-1.9166	-1.21	-.1412	2.0904	2.3467

Notes: Table shows descriptive statistics (without weights) on the explanatory variables for our applications. In particular, we use the share of OOP health expenditures in Section 5.2, while the percentage of inhabitants affected by an epidemic is used in Section 5.3. Finally, in Section 5.4, we use capital account liberalization as an explanatory variable for inequality.

Table D4: Regression Results Without Controlling for Country Area

	OOP Expenditure		Epidemics		CA Liberalization	
Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)
Income-based Gini	0.732*	1.442***	6.827**	5.761***	-0.186***	-0.210***
	(0.402)	(0.529)	(2.723)	(1.959)	(0.0431)	(0.0583)
Light-based Gini	0.990***	1.738***	6.433***	4.287**	-0.153***	-0.298***
	(0.280)	(0.509)	(1.637)	(2.136)	(0.0327)	(0.0773)
Observations	2089	2089	3278	3278	3013	3013
# of countries	177	177	187	187	177	188
Country fixed effects	no	no	no	no	no	no
Population weights	no	yes	no	yes	no	yes

Notes: Results without area controls of columns (1) and (4) of Table 2 for explanatory variable OOP expenditure, Table 3 for explanatory variable epidemics, and Table 4 for explanatory variable capital account liberalization. The light-based Gini-coefficient refers to the weighted measure with our preferred parameter choice of $\lambda = 0.5$. The income-based Gini-coefficient is obtained from Solt (2016). *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered at the country level and reported in parentheses.

Table D5: List of Countries and Territories

Country or territory	# inc	Country or territory	# inc	Country or territory	# inc	Country or territory	# inc	Country or territory	# inc
Afghanistan	6	Congo	7	Iceland	22	Myanmar	4	Slovakia	22
Aland Islands	0	Cook Islands	0	India	21	Namibia	21	Slovenia	22
Albania	17	Costa Rica	22	Indonesia	22	Nauru	1	Solomon Islands	9
Algeria	20	Cote d'Ivoire	22	Iran (Islamic Republic of)	22	Nepal	19	Somalia	1
American Samoa	0	Croatia	22	Iraq	7	Netherlands	22	South Africa	22
Andorra	0	Cuba	0	Ireland	22	New Caledonia	0	South Sudan	1
Angola	10	Curacao	0	Isle of Man	0	New Zealand	22	Spain	22
Antigua and Barbuda	1	Cyprus	22	Israel	22	Nicaragua	21	Sri Lanka	22
Argentina	22	Czech Republic	22	Italy	22	Niger	22	State of Palestine	16
Armenia	22	Denmark	22	Jamaica	13	Nigeria	19	Sudan	18
Aruba	0	Djibouti	18	Japan	22	Niue	0	Suriname	7
Australia	22	Dominica	9	Jersey	0	Norfolk Island	0	Swaziland	18
Austria	22	Dominican Republic	22	Jordan	22	Northern Mariana Islands	0	Sweden	22
Azerbaijan	17	DPR Korea	0	Kazakhstan	22	Norway	22	Switzerland	22
Bahamas	13	DR Congo	9	Kenya	16	Oman	0	Syrian Arab Republic	11
Bahrain	0	Ecuador	22	Kiribati	1	Pakistan	22	Taiwan	22
Bangladesh	22	Egypt	22	Kosovo	11	Palau	1	Tajikistan	22
Barbados	19	El Salvador	22	Kuwait	1	Panama	22	Thailand	22
Belarus	22	Equatorial Guinea	1	Kyrgyzstan	22	Papua New Guinea	14	Timor-Leste	13
Belgium	22	Eritrea	0	Lao PDR	22	Paraguay	22	Togo	9
Belize	17	Estonia	22	Latvia	22	Peru	22	Tokelau	0
Benin	12	Ethiopia	19	Lebanon	18	Philippines	22	Tonga	18
Bermuda	0	Faeroe Islands	0	Lesotho	19	Poland	22	Trinidad and Tobago	14
BES islands	0	Falkland Islands	0	Liberia	9	Portugal	22	Tunisia	21
Bhutan	10	Fiji	22	Libya	1	Puerto Rico	22	Turkey	22
Bolivia	22	Finland	22	Liechtenstein	0	Qatar	22	Turkmenistan	14
Bosnia and Herzegovina	13	France	22	Lithuania	22	Republic of Korea	22	Turks and Caicos Islands	1
Botswana	19	French Guiana	0	Luxembourg	22	Republic of Moldova	22	Tuvalu	17
Brazil	22	French Polynesia	0	Madagascar	21	Republic of North Macedonia	20	Uganda	22
British Virgin Islands	0	Gabon	1	Malawi	22	Reunion	0	Ukraine	22
Brunei Darussalam	0	Gambia	22	Malaysia	22	Romania	22	United Arab Emirates	1
Bulgaria	22	Georgia	22	Maldives	9	Russian Federation	22	United Kingdom	22
Burkina Faso	20	Germany	22	Mali	16	Rwanda	22	United Republic of Tanzania	22
Burundi	22	Ghana	22	Malta	15	Saint Helena	0	United States of America	22
Cambodia	16	Greece	22	Marshall Islands	0	Saint Kitts and Nevis	10	Uruguay	22
Cameroon	18	Greenland	0	Martinique	0	Saint Lucia	13	US Virgin Islands	0
Canada	22	Grenada	10	Mauritania	22	Saint Pierre and Miquelon	0	Uzbekistan	12
Cape Verde	15	Guadeloupe	0	Mauritius	21	Saint Vincent	14	Vanuatu	5
Cayman Islands	0	Guam	0	Mayotte	0	San Marino	0	Venezuela	22
Central African Republic	17	Guatemala	22	Mexico	22	Sao Tome and Principe	11	Viet Nam	22
Chad	9	Guernsey	0	Micronesia	16	Saudi Arabia	0	Wallis and Futuna Islands	0
Chile	22	Guinea	21	Monaco	0	Senegal	20	Western Sahara	0
China	22	Guinea-Bissau	18	Mongolia	19	Serbia	17	Western Samoa	7
China, Hong Kong	22	Guyana	16	Montenegro	9	Seychelles	8	Yemen	22
China, Macao	0	Haiti	12	Montserrat	0	Sierra Leone	20	Zambia	22
Colombia	22	Honduras	22	Morocco	22	Singapore	22	Zimbabwe	17
Comoros	10	Hungary	22	Mozambique	18	Sint Maarten (Dutch part)	0		

Notes: Table alphabetically lists all countries and territories in our sample and the corresponding number of observations for the income-based Gini-coefficient. For all 234 countries listed we have 22 yearly observations for the light-based Gini-coefficient and thus 5,148 observations in total.