



Measuring Inequality from Above

Jose G. Montalvo
Marta Reynal-Querol
Juan Carlos Muñoz-Mora

This versión: May 2021

Barcelona GSE Working Paper Series

Working Paper n° 1252

MEASURING INEQUALITY FROM ABOVE*

JOSE G. MONTALVO, UPF-BARCELONA GSE-IPEG-ICREA

MARTA REYNAL-QUEROL, ICREA-UPF-BARCELONA GSE-IPEG

JUAN CARLOS MUÑOZ-MORA, UNIVERSIDAD EAFIT

First version: January 2017

This version: May 2021

Abstract

Recent research has shown the usefulness of nighttime light (NTL) data as a proxy for growth and economic activity. This paper explores the potential of using luminosity at night, recorded by satellite imagery, to construct measures of inequality. We develop a new methodology to construct a Gini index for each country using the nighttime light per capita over millions of small pixels. To assess the usefulness of our procedure, we check the correlation of our measure with the common factor extracted from the analysis of several Gini indices calculated using traditional data sources. Finally, we show two specific applications of our methodology: the calculation of within and between inequality across regions and ethnic groups.

*Financial support from the European Research Council under the European Community ERC Grant Agreement n.647514, the Spanish National Science Foundation AEI/FEDER, UE ECO2017-82696-P, and the Government of Catalonia (ICREA) is gratefully acknowledged. We are also thankful for the financial support provided by the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2015-0563). In addition, the authors thankfully acknowledge the computer resources at MareNostrum and the technical support provided by Barcelona Supercomputing Center (FI-2018-1-0017, FI-2018-2-0044, FI-2018-3-008, FI-2019-1-0027, IM-2019-2-0013 and IM-2019-3-0016). J.C. Muñoz acknowledges the financial support under the project 'Inequality and Governance in Unstable Democracies: The Mediating Role of Trust', implemented by a consortium led by Institute of Development Studies (IDS). The support of the UK Economic and Social Research Council (ESRC) Grant: ES/S009965/1 is also kindly acknowledged.

1 Introduction

The sources of increasing inequality in recent times have become one of the most debated issues in economics. However, the measurement of inequality is a difficult topic even before moving into its interpretation. There are basically two types of issues. One question is the choice of a particular indicator to measure inequality from the basic source of income distribution. Popular measures include the Gini index, the Theil index or interdecile ratios. The second issue is how to choose the source of basic data. In addition to measurement errors in the scale factor, there may also be measurement errors associated with the shape factor, or the distribution derived from the income and consumption surveys. In particular, consumption estimates are sensitive to the length of the recall period and miss part of the consumption like the flow of consumption services. Income estimates are subject to nonrandom missingness, or no response, and under-representation of high-income earners. There are also many sources for income inequality depending on the data used and the specific measure that is being constructed. These issues are not only a difficult question for the measurement of inequality but also for the estimation of many other economic measures. For instance, different revisions of the Penn World Tables produce standard deviations of average growth rates that are almost as large as the average growth rates. It is also well known that, in the measurement of poverty, there are large differences between the results using the average income calculated using household income and expenditure surveys or using national account estimates.¹

Recently, in addition to National Accounts and income/consumption surveys, academic research has increasingly used satellite imagery to estimate economic activity. This approach has some advantages over the traditional indicators. The approach can be used at high levels of resolution (small areas) where it is difficult to find estimations of GDP, or enough individuals in national surveys to produce a reliable estimation. It is also a good approach to deal with the estimation of economic activity in war areas or places where there is a high level of social tension that make other approaches unfeasible. Finally, night light is measured in the same way around the world and, therefore, it simplifies the comparisons.

To measure income per capita in each cell, we need information on a measure of economic development and population. At high levels of resolution, it is difficult to find estimations of GDP and certainly, many areas of the world do not have information on geocoded high-resolution measures of economic development. It has, however, become increasingly common to use satellite nighttime light (NTL) density as a proxy for local economic activity when working with small geographical areas.

¹For instance Deaton (2005).

Satellite night-light density data are available from the National Oceanic and Atmospheric Administration. Previous research has shown that light density at night is a useful proxy for different measures. Sutton et al. (1997) found that the spatial analysis of saturated pixels predicted population with a R^2 of 0.63. Henderson et al. (2012b) finding an R^2 of 0.77 for the panel regression of growth in real GDP on the change in luminosity using year dummies. Chen and Nordhaus (2011) find that luminosity has informational value for countries, regions, and areas with poor quality or missing data. Using a traditional measurement error model, they argue that night light has a large estimated optimal weight in the estimation of growth rates in countries with low quality statistical systems, following the A to D classification of the Penn World Tables (PWT). In particular, Chen and Nordhaus (2011) show that the weight is, in these cases, larger than in the estimation of the level of GDP per capita. The importance of night light, as measured by its weight, in the estimation of growth is always higher in low-GDP density countries than in those of high-GDP density, for any level of quality of the statistical system.² More recently, Pinkovskiy and Sala-I-Martin (2016) have used nighttime lights to show that National Accounts are a good proxy for actual income, while income measured using survey means have very little, if any, informative content to estimate true income. They showed that growth rates of GDP per capita are very highly correlated with the growth of night light per capita while the growth rate of survey means is very weakly correlated with the growth of night light per capita. Jean et al. (2016) used satellite images and machine-learning techniques to predict poverty on small scales. In their application, they use daytime satellite photos to capture details of the landscape (metal roof, water, etc.) that they correlate, using neural networks, with satellite night lights as a proxy for economic activity. Other recent research using night light includes Michalopoulos and Papaioannou (2013, 2014), Alesina et al. (2016), Pinkovskiy and Sala-i Martin (2020) o Montalvo and Reynal-Querol (2020).

In this paper we use satellite imagery to calculate inequality measures instead of economic activity. As we argued in the previous paragraph, past literature has shown the usefulness of luminosity data as proxy for economic activity. Our objective is to show that these data can also be useful to measure economic inequality at high levels of geographical disaggregation and to decompose between and within groups inequality.

Thus, the idea is to find a procedure that can be generalised to a global scale and produce methodologically homogeneous measurements, particularly important when estimating inequality since there are significant differences in the methodology and calculation of the Gini index by different organisations that affect not only low income-low quality data countries but also middle-

²The cross-validation analysis in Michalopoulos and Papaioannou (2013) shows that light density at night is highly correlated with a wealth index across households in four large African countries.

and high-income countries. For instance, there are significant differences between the Gini indices calculated using LIS (Luxembourg Income Survey), which is an important reference for European countries, and the indices for those countries in the database of the World Bank. The average difference during the period 1988-2013 is close to 3 points. In the case of Spain, the National Statistical Office reported a Gini index of 33.7 in 2013 while the World Bank database reports an index of 36.2. These differences extend to the trend in the Gini coefficients in countries with a high index like US, South Africa or Brazil.

Our paper contributes to a growing literature that uses satellite imagery to proxy economic magnitudes. The first contribution is a new procedure to calculate inequality using luminosity and population measured using small pixels. The previous literature has relied on the value of nighttime light or difference in night light across regions or ethnic groups (spatial inequality). We weight economic activity by gridded population to proxy for income per capita. Second, it is well known that saturated night light, which is the one mostly used in the literature, is top-coded, and the use of different satellites can alter the measurement. We propose a procedure that allows us to estimate the optimal correction for the values of saturated nighttime light in areas where lights are very intense. In addition, instead of using the average of the nighttime light from all the satellites available in a particular period of time, we search for the one that generates the best fit. Once we have the corrected version of night-light we calculate the Gini index using the average light by pixel as the representative value for the individuals living in that pixel. We validate our inequality indicators using country data. After ranking all the pixels in each country we use the standard Gini formula to calculate the index based on nighttime light, which we denominate MIFA. We perform the validation analysis using country data since the alternative sources of Gini indices use country data. This approach makes the validation exercise quite challenging since countries are large units, and our indicator is expected to work better for small areas. Finally, we use our methodology to evaluate within and between regions and ethnic groups inequality.

The paper has the following structure. Section 2 locates the paper in the context of the literature on the measurement of inequalities using nighttime light, also introducing our methodological contribution in terms of the transformation of the data. Section 3 explains the sources of the data. Section 4 discusses a transformation to deal with the issue of the top coding of the saturated nighttime light measurements, and the choice of the parameters for such a transformation. Section 5 compares the results obtained using nighttime light with the inequality measures obtained using alternative sources of data. The purpose of this exercise is to obtain a common factor to all the available inequality indicators and to check the correlation between this common Gini factor and our measure. Section 6 calculates, using our methodology, inequality within and between regions

and ethnic groups across countries. Finally, Section 7 summarizes the conclusions.

2 Measuring Inequality From Above

The initial research on the use of nighttime lights as proxy for social indicators was devoted to its relationship with population density. Later research looked into the correlation between night light and economic activity. Only more recently, the attention has turned to other economic measures and, in particular, inequality. There are basically two approaches in the literature, although the basic view is the calculation of spatial inequality (regional, ethnic homelands, etc.). Lessmann and Seidel (2017) calculate regional predicted GDP per capita using several luminosity variables (average nighttime light within a region, top-coded pixels within the region and low-coded pixels within a region, among other variables). They use these predictions to calculate a Gini index that accounts for inequality across the regions of a country and claim that using predicted income generated inequality measures that are more correlated to the indices based on observed income than inequality measures simply based on nighttime light (NTL) density.

A second approach uses precise NTL dispersion to proxy for income inequality. The original contribution of Henderson et al. (2012b) was centred around the predictive ability of night light as a proxy of GDP per capita. In one of the tables, they show a Gini index calculated for eight countries using only NTL density³. Alesina et al. (2016) calculate ethnic inequality using the results of averaging luminosity of all the observations within the boundaries of an ethnic group and then dividing by population. Their Gini index is constructed using night light per capita for each ethnic group as one observation to be ranked in the construction of the Gini index.

Our main objective is computing inequality measures based on a proxy of income built from remote sensing data sources. Unlike traditional data sources, where individual census or tax data represent the income from a single person or household, our approach builds upon geographical units that represent a set of people living in a particular geographical area. Differently from other proposals using the night light referred to in the previous sections, that basically calculate between groups or spatial inequality, we calculate a measure of total inequality. Our approach tries to reproduce, as closely as possible, the calculation of a standard Gini index. The income of each individual is calculated as the average night light in a very small geographic unit (i.e., pixel) using data from satellite imagery combined with the corresponding gridded population. Many papers show that night light is a good proxy for economic activity, but it has some limitations related to

³Henderson et al. (2012b) Table 1. For the countries included in that table, there is a large difference between the Gini indices calculated using only light density and the standard Gini indices calculated using incomes surveys or national accounts.

the calculation of night light.⁴ In particular, research uses mostly the saturated version of night light which implies that the range of light is measured between 0 and 63. The presence of this top coding limits the ability of night light to capture high density of economic activity. The basic problem is that radiance light, which is not top-coded, has large measurement errors although radiance light is being found to be a better proxy for some measures.⁵ Instead of using radiance light, we use a correction for the saturated light that tries to overcome the problem of top-coding. We eliminate the pixels that do not have any associated population, which correspond mostly to 0 night-light scores. Thus, our income per capita proxy at the pixel level is defined as:

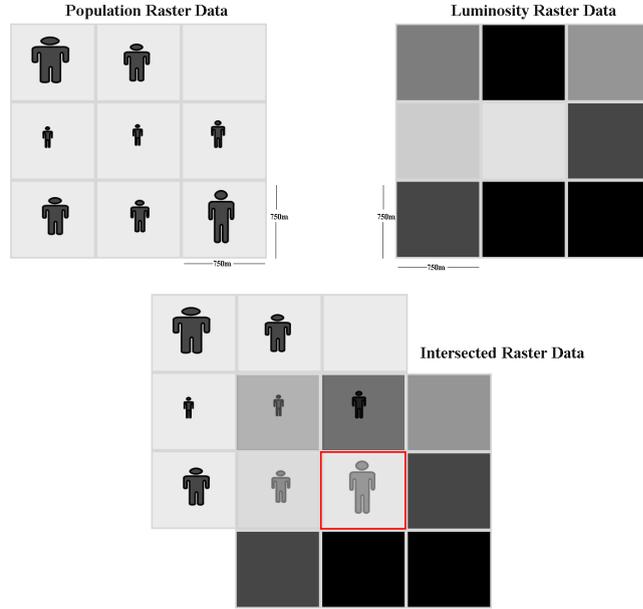
$$(1) \quad w_i = \frac{l_i}{p_i}$$

where l_i is our proxy of total income measured as the luminosity corrected, if that is the case, by top-coding in pixel i , and p_i is the number of people located at pixel i . That is, we use the luminosity per capita as the proxy of the level of economic development for a given geographic unit i . As a result, we create new raster data where each unit of analysis is a proxy of income per capita, calling this *income raster*. Elvidge et al. (2012) introduced a similar approach, *the Night Light Development Index -NLDI-*, which combines night-light and population satellite data to measure the inequality. For doing this, they aggregated the original pixel-level data into larger grids (10km x 10km), that might mask the problems such as top-coding and miss-aligned rasters. Figure 1 shows an example of building w_i using population and luminosity rasters with the same resolution (i.e., pixel size) and perfectly aligned. After overlapping both rasters, the new income raster is the result of intersecting both rasters in which each pixel contains the estimated w_i . In case of perfectly aligned rasters, the resolution of the income raster is the same as the input ones.

⁴We will discuss other limitations related to the obsolescence of satellites and its ability to correct for atmospheric effects.

⁵Mellander et al. (2015)

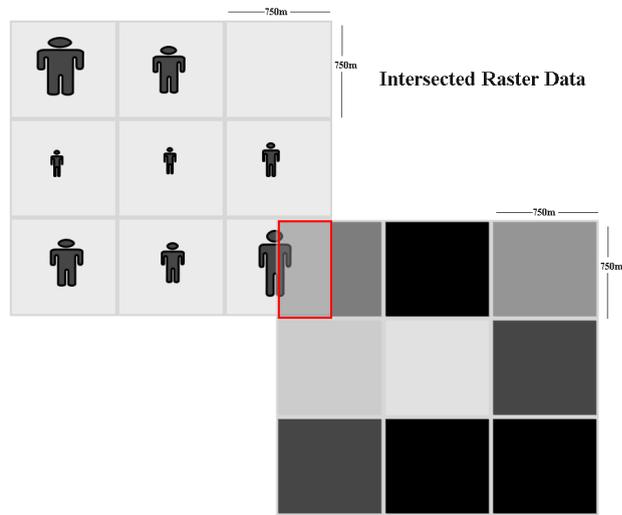
Figure 1: Example of building economic development approach at pixel level using a perfectly aligned rasters



Notes– This graph represents the main unit of analysis after the intersection of the population density and luminosity rasters with the same resolution and perfectly aligned.

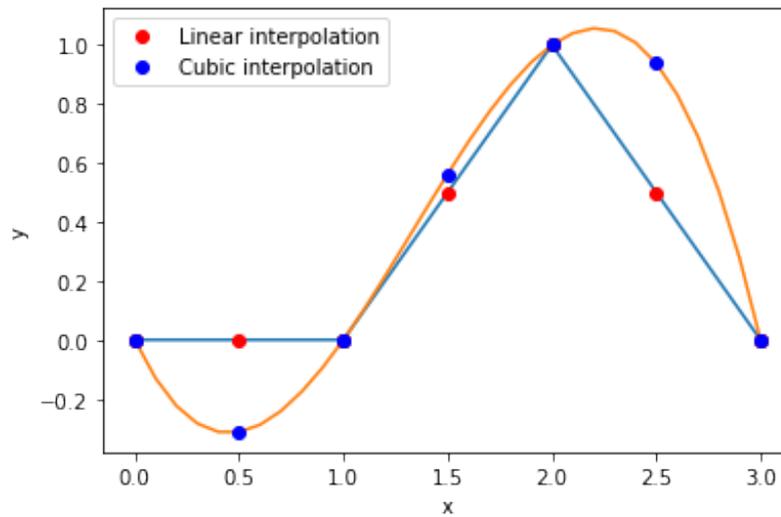
If we could access a perfectly aligned raster of luminosity and population, like the one shown in Figure 1, it would be simple to calculate our proxy of income per capita. However, this is not feasible since the basic information on both variables comes from different sources and methodologies. To construct the overlapping of the two rasters to produce the proxies w_i in a situation like the one in Figure 2, we need to make operational decisions. By assuming that both luminosity and population raster are homogeneously distributed within each pixel, it is possible to obtain l_i and p_i in the presence of nonaligned rasters. Since pixels have a very small size, this is not a very restrictive assumption. To assign a value for the new partitioned pixel resulting from intersecting the two nonaligned rasters, we interpolated its value depending on the original value of the pixels using interpolation methods. The two most used methods to perform this task are bicubic and bilinear interpolation. Even if the bicubic algorithm gives smoother shapes and better results in general, with abrupt population changes, negative values were obtained, notably on the coast. This is a well-known problem (bicubic interpolation can give values that are out of the bounds defined by the pixels in the frontier of the resampled one) (Keys, 1981). Hence, bilinear interpolation has been chosen. Figure 3 shows a 2D example to better understand the previously explained reasoning. In this case, the point 0.5 has a lower value regarding all the original input points.

Figure 2: Example of pixel level income per capita calculation using not aligned rasters



<https://www.overleaf.com/project/5d499792749eca36aa167dd6> *Notes- This graph represents the main unit of analysis after the intersection of the population density and luminosity rasters with the same resolution, yet not aligned.*

Figure 3: Example of linear and cubic interpolation in 2D



Notes- This graph represents the interpolated points each 0.5 units when using linear and cubic interpolations for the points (0,0), (1,0), (2,1) and (3,0).

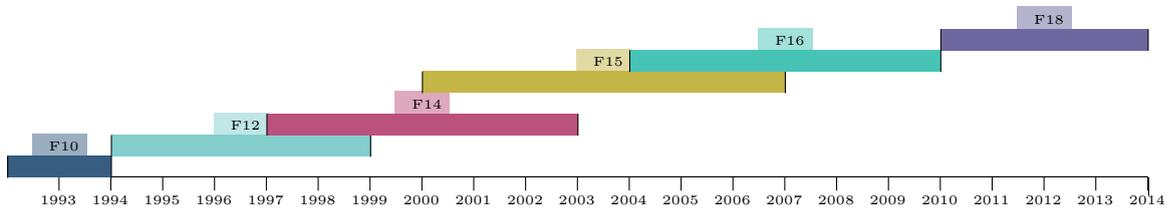
3 Data

The aim of this paper is to construct an index of inequality that uses a common methodology for all the countries and can cover any area of the world since it does not rely on the existence of surveys or National Accounts. The two basic inputs for our approach are luminosity and population.

3.1 Luminosity

We use the cloud-free night-light data collected by the Defense Meteorological Satellite Program - Operational Linescan System (DMPS-OLS) program⁶ processed by the National Geophysical Data Center at the National Oceanic Administration Agency -NOAA-, specifically the Earth Observation Group (EOG)⁷. EOG provides information on the average quantity of light observed at each pixel across cloud-free nights for every year at 30 arc-second resolution (approx. 750 m x 750 m), spanning from -180 to 180 degrees longitude and -65 to 75 degrees latitude⁸. Figure 4 shows the temporal availability for year and the different satellites. Its values range from 0 to 63, where 0 is the lowest luminosity and 63 the highest.⁹

Figure 4: Temporal availability by luminosity data set by satellite



Luminosity data have some potential sources of measurement error (Addison and Stewart, 2015). The first source of measurement error is associated to the high intensity of light in large cities. On the one hand, to capture this large spectrum of light intensity, a satellite would need a very sensitive-to-light lens and large storage capacity. Unfortunately, the data we use do not provide these requirements, capturing the intensity of light in truncated integer bins spanning from 0 to 63. As a consequence, the information available is a truncated distribution of intensity of light. That is, as information is top-coded at 63, and it is not possible to distinguish the difference between the main downtown and suburbs in big cities as they are coded as 63. This problem is known as top-coding¹⁰. On the other hand, even if the lens were sensitive enough to capture the entire spectrum of light, there are physical phenomena of light (e.g., light reflection, world

⁶The nighttime light data from the DMPS-OLS has been a primary data source for most of the previous economic literature.

⁷For further information and data at <http://www.ngdc.noaa.gov/eog/>. NOAA also generates series of radiance-calibrated lights combining the stable lights data with auxiliary information from low amplifications sensors.

⁸In the appendix there is a detailed description of each dataset

⁹The information of the radiance-calibrated lights are not top-coded and, therefore, there is no upper-bound in the value. However, they are very limited for time comparison, they are quite unstable and infrequently measured. Recently, the Suomi National Polar-orbiting Partnership (Suomi-NPP) satellite using the Visible Infrared Imaging Radiometer Suite (VIIRS) has begun producing information on nighttime light. In principle, the radiometric and spatial resolution of the VIIRS has improved with respect to the DMPS. It has few over-glow effects and it is more precise with light within cities. However, at the time we began this research the information was not available, and it is only produced since 2013.

¹⁰To see some examples of the top-coding problem for the main cities see Rosenbloom (2009)

curvature, among others) that might over-magnify the size of the city. In some cases, cities appear magnified to approximately 10 times their true size (Imhoff et al., 1997; Henderson et al., 2003; Small et al., 2005). This problem is known for creating a problem known as pervasive blurring, dubbed 'overflow' or 'blooming'¹¹.

In contrast with large cities, the second source of measurement error is in those places where the intensity of light is so weak that it is very similar to the background illumination of the earth. Therefore, satellites often code places with very weak light either as zero light or very low values (below 5). The final and third source of potential bias is the lack of comparability in absolute term of the data. As satellite technology has been improving over time with more sensitive lenses, large storage capacity, among other improvements, there is a lack of comparability across years with different satellites. For solving this problem, Hsu et al. (2015) and Wu et al. (2013) provides adjustment parameters to allow temporal comparability across satellites¹².

In our case, we would like to estimate relative measures of inequalities based on building a new wealth raster. The main sources of measurement error are those associated with the behaviour of the night-light intensity at both ends of the distribution (i.e., top-coding and zero light). We have devoted a extensive work to minimise their potential effect over our main results. The exact procedure will be described in the following sections.

3.2 Population

To capture the geographical population distribution at very high resolution, we used the Landscan Global model, which represents an 'ambient' population distribution over a 24-hour period by integrating diurnal movements and collective travel habits into a single measure (Dobson et al., 2000). The model combines a multilayered, dasymmetric, spatial modeling approach (also known as "smart interpolation"), to reallocate population within the lowest geographical units available in each country (i.e., census units) (Dobson et al., 2000). This modelling process is based on the census units for each country and primary geospatial input or ancillary data sets, including land cover, roads, slope, urban areas, village locations, and high resolution imagery analysis (Lanscan, 2018). As result, the process offers raster data representing human population distributions at 30 arc-second resolution (approx. 750 m x 750 m) worldwide from 2000 until 2019. These data are widely used for risk management and policy intervention around the world (Dobson et al., 2000).¹³

¹¹Abrahams et al. (2018) provide a very comprehensive review of the blurring bias for the main cities around the world.

¹²Another problem associated with night-light data is light emissions that are not associated with an economic activity such as gas flares. In these cases, NOAA provides gas flare boundaries, which allow exclusion of those specific pixels, solving the problem

¹³Appendix A offers a review on census-based alternatives to decompose census units into smaller geographical units.

3.3 Inequality index

The income raster developed in the previous section provides a geographically explicit continuous distribution of nighttime light at a very low geographical level available worldwide. This information is the main input to calculate inequality indices like the Gini index. The Gini index is a measure of inequality, defined as the mean of absolute differences between all pairs of individuals within a given unit of analysis (e.g., administrative borders, social groups, among others). Thus, it can be interpreted as the expected income gap between two individuals randomly selected from the population (Sen and Foster, 2005). There are several alternatives to estimate the Gini index (Yitzhaki and Schechtman, 2013). The final choice among the different alternatives is based mainly on the types of distribution (e.g., continuous, nonnegative values, among others) or the use of weights. The classical notation for Gini index based on the theory of relative mean difference is given by:

$$(2) \quad \text{Gini} = \frac{\sum_{i=1}^n \sum_{j=1}^n |w_i - w_j|}{2n^2 \frac{\sum_{i=1}^n |w_i|}{n}} = \frac{\sum_{i=1}^n \sum_{j=1}^n |w_i - w_j|}{2n \sum_{i=1}^n w_i}$$

where w_i stands for the nighttime light per pixels and n is the number of observations.

To calculate the Gini index, we rank the light per capita of each pixel from the lowest to the highest values. Table 1 shows the basic statistics derived from the methodology proposed in the previous section. We consider 186 countries. Our methodology generates an average of 4.41 million of pixels per country. However, the average number of populated pixels by country is 2.22 million.¹⁴ The average of the proportion of populated pixels over total pixels across countries is 85%. This implies that there is a large dispersion across countries in the proportion of populated pixels.¹⁵ Canada and Iceland have proportions below 25% while the proportion is 100% in San Marino, Monaco or Malta. The country average of inhabitants per populated pixel is 81.6.

Table 1: Summary statistics

Variable	Number of countries	Mean	St. Dev.
Num. pixels by country (millions)	186	4.41	15.3
Num. of populated pixels by country (millions)	186	2.22	4.98
Percentage of populated pixels	186	0.85	0.21
Population by populated pixel	186	81.6	510.6
Percentage of pixels DS > 50	186	0.048	0.14
Percentage of top-coded pixels DS = 63	186	0.016	0.09

The literature has recognised three basic issues when working with luminosity as a proxy for income. The first one is the censoring of data derived from the top-coding of very high luminosity

¹⁴We exclude from our calculation the pixels that have no population.

¹⁵The proportion of populated pixels over the total number of pixels is close to 50%

pixels in the saturated version. In our data 1.6% of the pixels are top coded and 4.8% are above 50. The second problems is the low, or zero, coding of some pixels. This is not problematic if there is no population. For pixels that have population we set the value to the minimum luminosity per capita of the country.¹⁶ Finally, the sensitivity of the light measurement depends on changes in satellites and sensor technology. The following section discusses these issues.

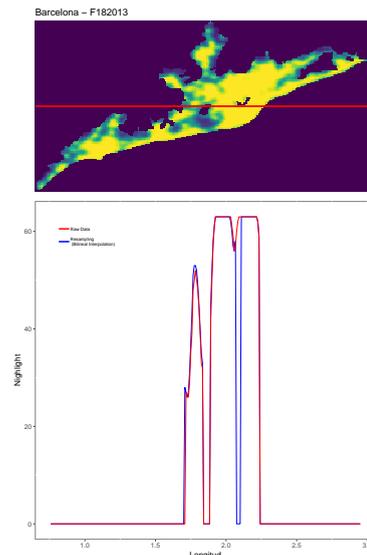
4 Measuring inequality using luminosity

In this section, we describe the methodology to deal with the issues that we discussed in the description of the luminosity data and the alignment of the rasters.

4.1 Interpolation

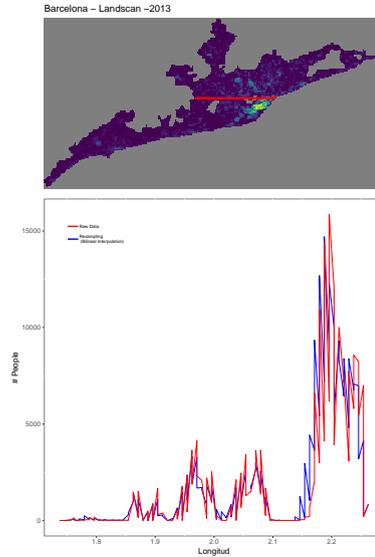
In the previous section, we discussed the interpolation procedure. In this section, we present the application of the procedures to our data with some examples to clarify the method. The first stage is to interpolate the data to align rasters of income and population. As an example, we show the application of the interpolation to the case of the area around the city of Barcelona. Figure 5 shows the interpolation of luminosity and Figure 6 represents the interpolation of population.

Figure 5: Bipolar Resampling -Year 2013 - Satellite F18



¹⁶In the applications in Section 6 we use the minimum value of luminosity per capita of the region/ethnic group that the pixel belongs to.

Figure 6: Bipolar Resampling -Year 2013 - Population



4.2 Dealing with the top-coding at pixel level: a local Pareto correction

As we pointed out, in our basic data source nighttime light (NTL) is upper-bounded at 63, leading to some urban areas with uniform values. Initially, scholars used to disregard this problem by creating a threshold above which data would be discarded. This approach was rapidly abandoned, and in many cross-country analyses using night-light, this problem is not even mentioned. Elvidge et al. (2009) and Hsu et al. (2015) developed an alternative solution by triangulating the original data with complementary sources. They proposed the combination of the DMSP-OLS data set with auxiliary data set obtained from the pre-flight sensor calibration. As a result, it is possible to have a radiance calibrated NTL free of top-coding. The OLS radiance calibrated NTL might potentially solve the saturation bias, there are some important caveats. First, as the pre-flights calibration are seldom done, the data is only available in seven years (1996,1999,2000,2003,2004,2006 and 2010). Second, because the DMSP-OLS does not have any on-board calibration device, there is not a clear baseline value from the actual level of saturation. Thus, the radiance values should still be considered relative and not absolute (Hsu et al., 2015). Third, there are important differences in the time frequency with year (fewer orbits) and geographical coverage for the calibration values might varies across time. As a result, the radiance calibrated products present high instability in its quality and coverage (Hsu et al., 2015).

An alternative solution to overcome the top-coding problem is to assume a particular distribution for nighttime light. In urban economics the upper tail of the size of the distributions has been

model as a power law distribution.¹⁷ Small et al. (2011) analyze the urban extent of cities using night lights and shows that the size distribution of big agglomerations follows a Pareto distribution, a particular power law distribution. This distribution is also a common choice for the distribution of top income¹⁸. In addition Bluhm and Krause (2018) argue that the Pareto distribution is a good description of top lights. Following these insides, we are going to assume that the distribution of light above the threshold follows the Pareto law. That is, the top-light distribution could be described following a cumulative distribution function $F(l)$ for luminosity l , defined by:

$$(3) \quad 1 - F(l) = \left(\frac{l_m}{l}\right)^\alpha \quad (l_m > 0, \alpha > 1)$$

where l_m and α are exogenous values, also known as Pareto parameters. The corresponding density function could be expressed as $f(l) = \alpha \frac{l_m^\alpha}{l^{1+\alpha}}$. The assumption that the actual luminosity distribution follows a Pareto distribution implies that the ratio of average light intensity at the top pixels over the threshold income does not depend on the level of the threshold.

$$(4) \quad \frac{l^{top}}{l} = \beta \quad \text{with} \quad \beta = \frac{\alpha}{\alpha - 1}$$

where β is known as the inverted Pareto coefficient, which intuitively indicates that higher β means a fatter upper tail of the distribution (Atkinson et al., 2011).

The Pareto correction assumes that each observation is independent. We name this version the static neighbourhood approach. That is, the location of each observation is not considered when interpolating the value of the top-coded observation. However, using luminosity data, it is difficult to support this assumption, as data have an intrinsic spatial correlation. That is, pixels around a large city will be more likely to be top-coded than pixels located in rural areas. Likewise, the magnitude of the top-coding will depend on its location relative to their neighbourhood, so even if a pixel is not over the top-coding threshold, its value might be overmagnified (e.g., blurring effect of light). As a consequence, the Pareto correction needs to explicitly include the spatial location when interpolating the top-coded values.

To overcome this issue, we built a local version of the Pareto correction in which each top-coded observation was corrected relative to the distribution of non-coded observations located

¹⁷Rozenfeld et al. (2011)

¹⁸For instance Atkinson et al. (2011)

in its neighbourhood. That is, once a top-coded pixel was detected, we defined a $N \times N$ square neighbourhood in which the studied pixel would be the centre, then we used this distribution to correct the top-coded pixel. As we pointed out above, it is very likely that top-coded pixels are clustered around themselves, so it is possible that within a $N \times N$ neighbourhood of a given top-coded pixel, we can have several top-coded pixels. In those cases, all the values that contain a top value in their neighbourhood will be modified. Nevertheless, this change will be conditioned by the number of top level pixels in it.¹⁹

Therefore, we can define the local Pareto correction as follows:

$$(5) \quad l_{corrected} = \frac{\sum_{l_A \in A} l_A + \frac{\alpha}{1-\alpha} \sum_{l_B \in B} l_B}{N \cdot N}$$

where N is an odd number that indicates the size of the square defining the neighbourhood. A is the set of non-top-coded pixels located in the $N \times N$ neighbourhood, and B are the sub-set of the neighbourhood top-coded pixels. We name this version the dynamic neighbourhood approach. Finally, α is the Pareto parameter. Figure 7 shows the application of this methodology to the city of Barcelona assuming an $\alpha = 2$. Both the static and the dynamic neighbourhood techniques give similar results. Nevertheless, the second technique is smoother than the first one.

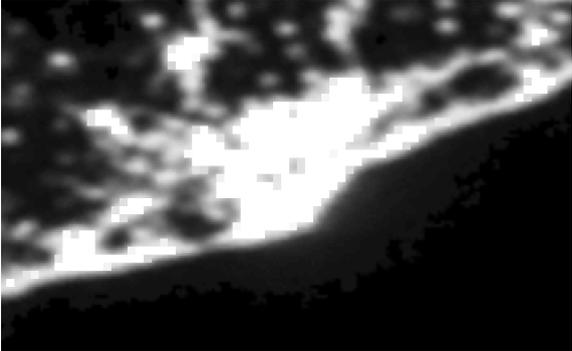
¹⁹Bluhm and Krause (2018) provide a top-coded correction using the Pareto distribution. However, in this case, all the top-coded pixels in a given region are considered when performing the corrections. Thus, this change will be conditioned by the number of top-level pixels in it. Indeed, the new value will be the average of the non-top-coded values and the corrected top-coded values present in the neighbourhood considered. Hence, the more top-coded values in the considered square, the higher the change. In addition, a global study on each top-coded concentration is done to sort the pixels considering their distribution, assuming a monocentric and circular shape. This sorting makes the method computationally intensive and is not viable when considering the whole world. We have decided to consider only the near pixels when correcting their value.

Figure 7: Changes in the rasters with the explained techniques in Barcelona's area

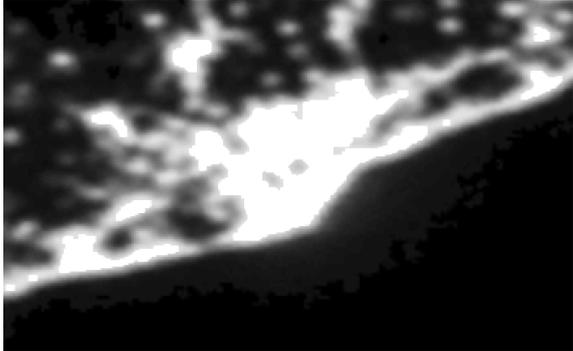
Original raster



Static neighbourhood



Variable neighbourhood



To choose the parameter for the local Pareto correction, we calculate the optimal value over different values of α . In particular, we create a local Pareto correction for the entire world for $\alpha \in \{1.25, 1.5, 2, 2.5, 5, 50\}$, or equivalently, $\beta \in \{5, 3, 2, 1.6, 1.2, 1\}$. Likewise, after considering different values, we set the neighbourhood size at $N = 5$.

Once the luminosity and population data sources have been correctly treated following the procedures explained above, the new light per capita raster maps are generated.

4.3 Choosing the optimal α

To complete the correction of the luminosity data, we need to fine-tune our procedure. There are two basic parameters that we need to calculate: (i) the Pareto parameters (α) for the top-coding correction; and, (ii) the optimal satellite for years in which there are data from more than one satellite available²⁰. The objective is to choose the optimal α comparing the difference between measures of inequality calculated not using satellite imagery and our estimate using a particular parameter.

We approach this selection heuristically, choosing first the α parameter from the set we described before and, afterward, choosing the optimal satellite for each year. Therefore, we calculate the α to minimise the following expression²¹:

$$(6) \quad \text{Min}_{\alpha} \sum_{k=1}^K \sum_{i=1}^N \sum_{t=1}^T (Gini_{i,t,k} - Gini_{i,t,\alpha}^{MIFA})^2$$

where i refers to countries, t to years, k to the alternative Gini indices exposed below, and α is the parameter for the correction of the top-coded pixels for our MIFA index.²²

For comparison purposes we use alternative sources of inequality measures.²³ There are several international organisations and research centres that compile datasets by country. They differ in coverage, data sources and indicators. For the purposes of this research we are going to consider only Gini indices. The basic difference between these alternative Gini indices is the source of data, which also determines the temporal coverage. Some of the indices are calculated from primary sources while others use secondary sources. Among the indicators that use primary sources, there are two types: those that try to be consistent with the definitions of National Accounts, and those based on microdata from surveys.

²⁰Previous research has mostly chosen to calculate the average luminosity across the available satellites.

²¹All the calculation were performed using the MareNostrum 4 supercomputer at Barcelona Supercomputing Center.

²²In principle, we take the average across the Gini indices calculated with all the available satellites. Later, we discuss the selection of one satellite for each year.

²³See the Appendix for more information.

A well-known problem of the calculation of inequality using microdata is that they severely underrepresent the extent of high income individuals. A notorious example of the data that tries to correct that problem using National Accounts is the World Income Database (WID) described by Piketty and Saez (2014). The differences in the type of data used to calculate the inequality indices lead to inconsistencies both in the levels and the trends. In addition, there is a clear tradeoff between coverage and comparability: increasing coverage means decreasing quality (imputation of missing values) and comparability (using sources that measure inequality using different primary sources).

For the validation exercise, we considered two indices constructed from primary sources and two indices that use secondary sources. Table 2 summarizes the data sources.²⁴

Table 2: Country-Level Income Inequality Data

Data Set	Host Institution	Source data
PovcalNet/World Development Indicators (WDI)	World Bank	National household surveys and Luxembourg Income Study (LIS). household survey statistics obtained from national statistical offices, the SocioEconomic Database for Latin America and the Caribbean (SEDLAC), the OECD Income Distribution database (IDD), the EU-Statistics on Income and Living Conditions (EU-SILC), LIS and PovCalNet.
World Income Inequality Database (WIID)	UNU-Wider	Fiscal (income tax) data and data from national accounts combined with other sources (household income and wealth surveys, inheritance and wealth tax data, as well as wealth rankings published in the media).
World Wealth and Income Inequality Database (WID)	World Inequality Lab	The project builds on various existing data resources, including the LIS.
Global Consumption and Income Project (GCIP)	Arjun Jayadev, Rahul Lahoti, Sanjay G. Reddy	

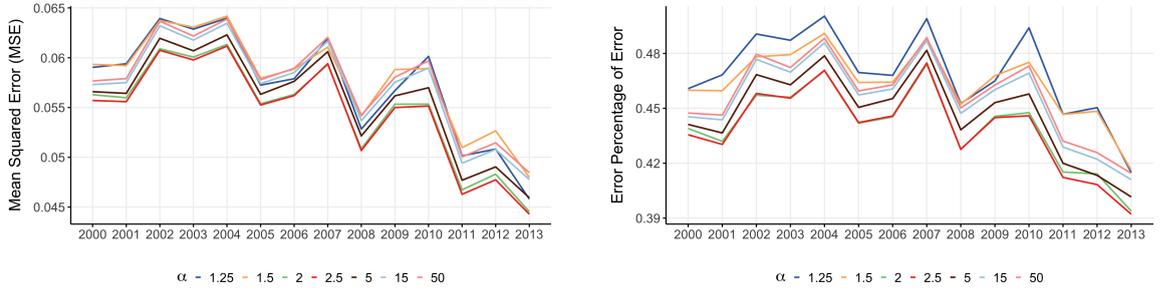
In the first group we include the estimates of the World Bank (WB) and the World Inequality Dataset (WID). The World Bank uses more than one thousand household surveys to construct their index while the WID uses fiscal information and National Accounts.²⁵ In addition, we consider the two largest datasets constructed using secondary sources. These datasets have, logically, a larger coverage than the previous datasets, but they include less comparable data. The World Income Inequality Database (WIID) includes information on 182 countries from many different sources (including the World Bank). The Global Consumption and Income Project (GCIP) presents the largest coverage of all the secondary datasets available.

²⁴The Appendix describes in detail each of the sources and the construction of the indices. We have chosen databases that are currently updated or that have a very large coverage.

²⁵Among the different versions of the WID data, we take the one that has the largest coverage, which is three times the coverage of the other versions.

Overall, there are trade-offs between coverage and comparability when developing income and inequality data sets. On the one hand, to develop a comparable and homogenised inequality measure, researchers might require harmonised data sources that guarantee same methodologies and time coverage. These are extremely difficult even among developed countries. Thus, maximising comparability and quality means focusing on a small number of countries, often developed countries. On the other hand, increasing coverage might require several data sources with different variables and data methodologies to produce estimates or performance imputations. This fact reduces comparability and quality of data sources.

Figure 8: Choosing the optimal top-coding correction α



Notes— This graph presents the average Mean Square Error (panel a) and Percentage of Error (PE) (panel b) from 2000 until 2013 across the seven data sources for eight alphas.

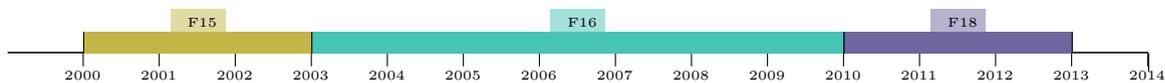
Figure 8 shows the average MSE from 2000 until 2013 across the four data sources. We use eight different Pareto correction values, $\alpha \in \{1.25, 1.5, 2, 2.5, 5, 15, 50\}$. On average the MSE for all α 's and years is 0.055. The mean squared error is reduced over time for all the possible values of α . The best performance was obtained using $\alpha = 2.5$ ²⁶. Setting $\alpha = 2.5$, we perform the similar analysis for the different satellites available in each year, where $sat \in (F14, F15, F16, F18)$.

$$(7) \quad \text{Min}_{sat} \sum_{k=1}^K \sum_{i=1}^N \sum_{t \in sat} (Gini_{i,t,k} - Gini_{i,t,\alpha=2.5}^{MIFA})^2$$

Figure 9 presents the satellite-year that minimises the mean squared error. Not surprisingly, the newest satellite in each period is the one that fits better.

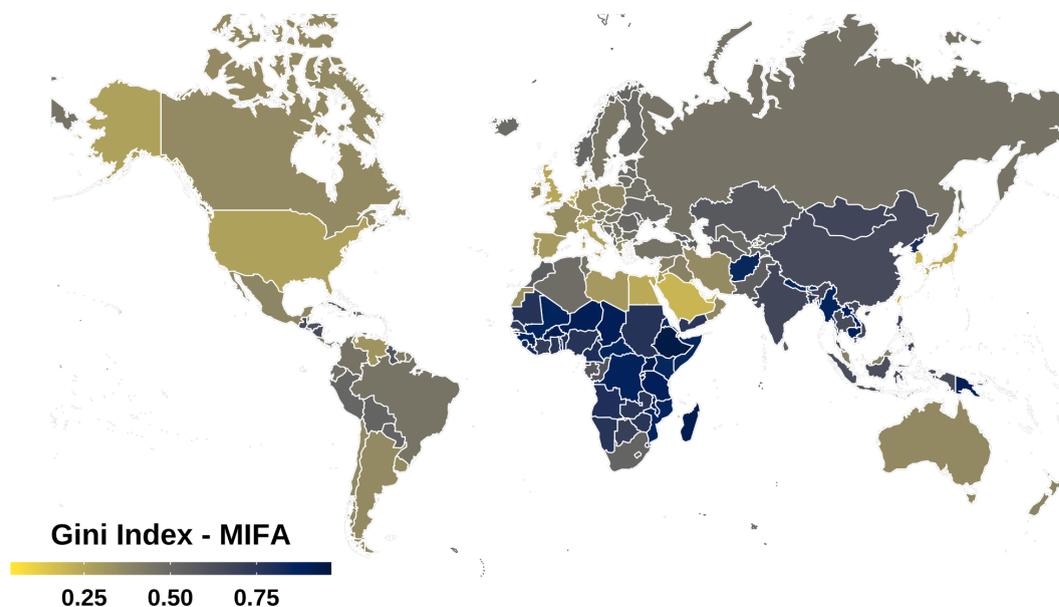
²⁶This parameter implies that the average value above the threshold is 67% higher than the value at the threshold. We should notice that the optimal α could also be calculated letting its value differ across countries. We follow the more restrictive approach of assuming a common parameter for all the countries.

Figure 9: Satellite that minimise the MSE each year



Using the optimal choices of α and satellite our Gini index produces the Figure 10 that shows the spatial distribution of the average GINI MIFA from 2000 to 2013. The highest indices are found in Africa where our estimates are quite high. In most developed countries, we find indexes that the value varies from 0.3 to 0.5.

Figure 10: Average Gini from Above



5 Extracting the common Gini factor

The objective in this section is to calculate inequality as a common factor among the five Gini indices that we have described in the data section: the World Bank (WB), the Global Consumption and Income Project (GCIP), the World Income Inequality Database (WIID) from the UNU-WIDER, the World Inequality Database (WID) and our indicator for Measuring Inequality from Above (MIFA). This approach implies that if there are missing values in any of the non-MIFA indicators, that observation is not considered in the common factors estimation. Table 1 presents the basic statistics of the observations of the five Gini indices included in the balanced sample.

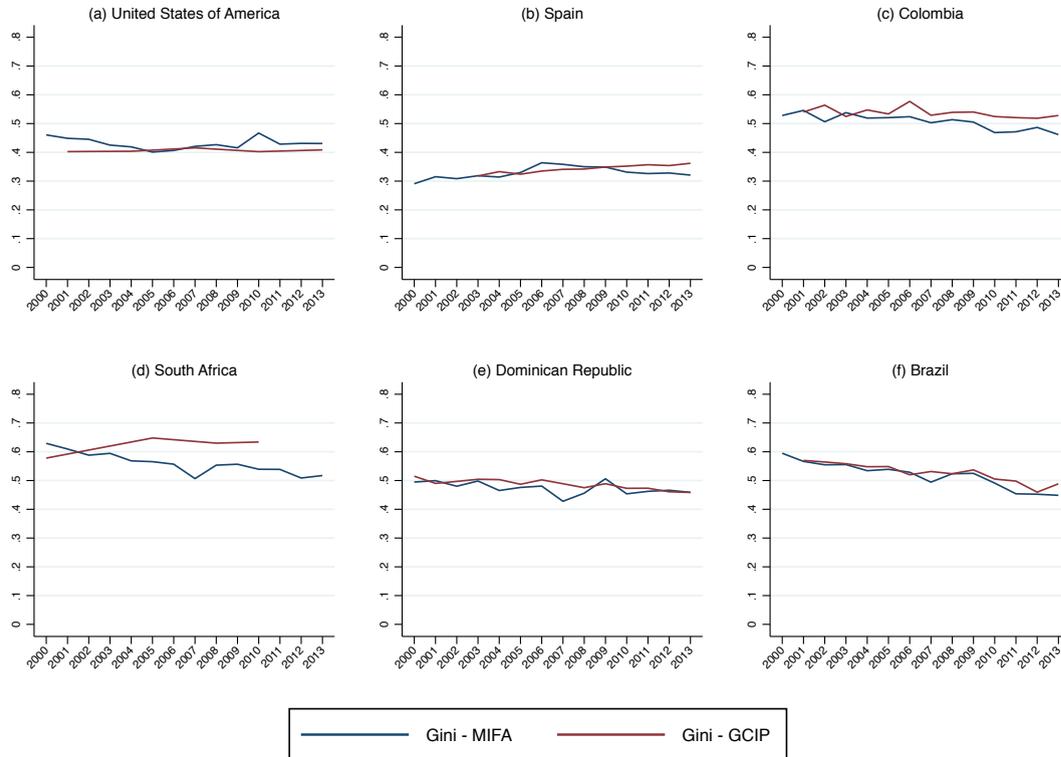
The mean of the indicators ranges from 0.35 to 0.48. The highest average Gini index corresponds to the WID indicator. The largest standard deviation corresponds to the MIFA indicator.

Table 3: Summary Statistics, 2000-2013

Statistic	Mean	St. Dev.	Min	Max
Gini - WB	0.351	0.075	0.237	0.647
Gini - GCIP	0.398	0.119	0.236	0.852
Gini - WIID	0.354	0.075	0.239	0.661
Gini - WID	0.487	0.091	0.331	0.776
Gini - MIFA	0.461	0.189	0.128	0.950

In addition to the difference across countries, we can also look at time variability within each country, especially among the countries with high levels of inequality. Figure 11 shows the evolution over time of the Gini index of some representative countries comparing the MIFA index with the Gini in the largest dataset among the other sources, which is the GCIP. The similarity of the evolution of the MIFA Gini and the GCIP index is quite remarkable in countries like Brazil, Dominican Republic, Spain and the US. In the case of South Africa and Colombia, the correspondence between the two indices is less remarkable.

Figure 11: Average Gini from Above

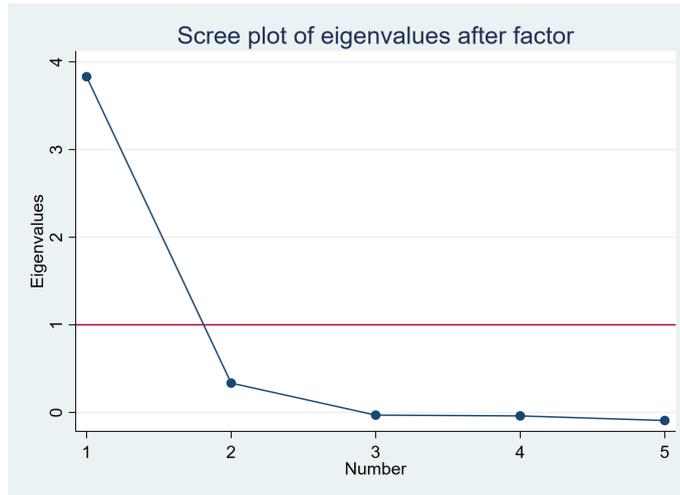


To evaluate our index we need to compare it with the unobserved common factor, or factors, which drive the co-movements of the different indicators of inequality discussed previously. First of all, in this section, we run an exploratory analysis to determine the number of factors associated with all these indicators. In a second stage, we specified an SEM model to estimate the unobserved factor. Finally, we analyse the determinants of the difference between our proposed Gini index and the estimated factor.

5.1 Exploratory factor analysis

The first stage is to explore the number of factors that correspond with the indicators. In principle, since all the indicators are measuring inequality, there should be a single common factor. The scree plot is the standard instrument to determine the number of common factors in our setup. We obtain the scree plot, or representation of the eigenvalues, presented in Figure 12.

Figure 12: Scree plot



The scree plot shows that there is only a very strong common factor with an eigenvalue clearly above 1. The other factors have very low eigenvalues. In addition, the overall Kaiser-Meyer-Olkin (KMO) is 0.84 which implies that the overall sampling adequacy is appropriate. The KMO for each indicator shows that they are also adequate. The three highest KMOs are associated with the WID indicator (0.88), the GCIP indicator (0.86) and the MIFA indicator (0.84). Therefore, the data seem appropriate for factor analysis, and there is one common factor to all these indicators.

5.2 Confirmatory factor analysis

The second stage is to characterise the relationship between the indicators, the common factor and the measurement errors. Following the results of the exploratory analysis, the specification of our

problem corresponds to a one-factor model. This framework is similar to the proposal in Pinkovskiy and Sala-i Martin (2016), Henderson et al. (2012a) and Chen and Nordhaus (2011). However, our approach to this signal extraction problem is different from the approach in this previous research. We assume the classical set up of a factor model: the Gini indicators are a linear function of the true Gini, measured with some error,

$$\begin{aligned}
(8) \quad G_{i,t}^{MIFA} &= \lambda_{MIFA} G_{i,t}^* + \epsilon_{i,t}^{MIFA} \\
G_{i,t}^{WB} &= \lambda_{WB} G_{i,t}^* + \epsilon_{i,t}^{WB} \\
G_{i,t}^{WID} &= \lambda_{WID} G_{i,t}^* + \epsilon_{i,t}^{WID} \\
G_{i,t}^{WIID} &= \lambda_{WIID} G_{i,t}^* + \epsilon_{i,t}^{WIID} \\
G_{i,t}^{GCIP} &= \lambda_{GCIP} G_{i,t}^* + \epsilon_{i,t}^{GCIP}
\end{aligned}$$

where $G_{i,t}^*$ is the true Gini and the ϵ s are the error terms which are uncorrelated with $G_{i,t}^*$

$$(9) \quad E(\epsilon_{i,t}^k G_{i,t}^*) = 0$$

Therefore, we can represent the system, including constants, as

$$(10) \quad X - \mu = \Lambda f + \epsilon$$

where Λ is the vector of factor loadings and f is the common factor. We assume that $E(f\epsilon) = 0$ and $E(\epsilon) = 0$. However, and opposite to the traditional orthogonal factor model, we cannot assume that $E(\epsilon\epsilon') = \Psi$ where Ψ is diagonal. Following the previous literature, we assume that the error term of the measurement constructed using night-light, in our case the Gini-MIFA, is uncorrelated with the other indicators. We also assume that the WID is uncorrelated with other indicators. However, the GCIP and the WIID are secondary indicators that use the WB and LIS as some of their sources.²⁷ Therefore, there is potentially a correlation between the measurement error in the WB indicator and the errors in the WIID and the GCIP indices. Logically, this correlation implies that there is potentially a correlation between the errors of the GCIP and the WIID indicators.

²⁷The algorithms to choose among the different datasets available are different for the GCIP and the WIID indices.

Considering those assumptions the problem we need to solve is

$$(11) \quad \Sigma = \Lambda\Lambda' + \Psi$$

where Σ is the covariance matrix of the Gini indicators. Using the assumptions described before, we have an overidentified system since there are 15 equations and 13 parameters to be estimated. Table 4 shows the results of the estimation of the system. To identify the variance of the latent Gini, we need to fix either a load or the variance of the latent Gini. The first column includes the results of the estimation using a fixed load for the first indicator, which is called the reference indicator. The second column presents the results when fixing the variance of the latent variable to 1. Since the variance of the latent Gini has been fixed to 1, there is no need for a reference indicator. Both columns show that all the loads and the variances of the errors are statistically significant.²⁸ The covariances are also significant with the exception of the one between the World Bank Gini and the GCIP. The final column presents the estimates of the standardised loadings, which is helpful for the interpretation of the results. For instance, one standard deviation on the latent Gini indicator will imply an increase of 0.59 standard deviations of the MIFA index, 0.91 standard deviations of the World Bank index or 0.98 standard deviations of the WID indicator.

To assess the goodness of fit of the model, there are several alternatives. How closely is the one-factor model to fit the data? The Comparative Fit Index (CFI) is generally used to measure this fit by comparing the estimated model with a baseline model that assumes no relationship between our five observed indicator variables. It is calculated as 1 minus the ratio of the chi-square of the estimated model divided by the chi-square of the baseline specification. The CFI of the estimated structure is 0.991, which indicates that the model is appropriate, and the standard cutoff for the CFI is approximately 0.9. Finally, the root mean squared error of the approximation is approximately 0.055, which shows that the fit is reasonably good. Finally, the reliability of this type of model is calculated as

$$(12) \quad \rho = \frac{(\sum \lambda_k)^2}{(\sum \lambda_k)^2 + \sum \sigma_{kk}^2 + 2 \sum \sum \sigma_{kk'}}$$

where k represents each of the indicators, λ_k are the unstandardised loadings, σ_{kk}^2 is the variance of the measurement error of each indicator, and $\sigma_{kk'}$ is the unstandardised covariance of the errors. The reliability is 0.79, which implies that the true unobserved Gini can account for almost 80% of

²⁸Results are robust for using other estimation procedures and the calculation of the standard errors of the coefficients using a robust, or bootstrap, estimator.

Table 4: Structural Equation Models

	Unstarndardized (1)	Gini = 1 (2)	Standardized (3)
MIFA			
Gini	1 (constrained)	0.111*** (0.008)	0.590*** (0.030)
Constant	0.461*** (0.008)	0.461*** (0.008)	2.443*** (0.088)
WB			
Gini	0.612*** (0.040)	0.0682*** (0.003)	0.913*** (0.009)
Constant	0.351*** (0.003)	0.351*** (0.003)	4.696*** (0.152)
GCIP			
Gini	0.938*** (0.052)	0.104*** (0.004)	0.878*** (0.017)
Constant	0.398*** (0.005)	0.398*** (0.005)	3.342*** (0.113)
WIID			
Gini	0.598*** (0.039)	0.0667*** (0.003)	0.885*** (0.011)
Constant	0.354*** (0.003)	0.354*** (0.003)	4.704*** (0.153)
WID			
Gini	0.810*** (0.050)	0.0903*** (0.003)	0.988*** (0.007)
Constant	0.487*** (0.004)	0.487*** (0.004)	5.337*** (0.172)
Var(ϵ^{MIFA})	0.0232*** (0.001)	0.0232*** (0.001)	0.651*** (0.034)
Var(ϵ^{WB})	0.000929*** (0.000)	0.000929*** (0.000)	0.167*** (0.017)
Var(ϵ^{GCIP})	0.00324*** (0.000)	0.00324*** (0.000)	0.229*** (0.021)
Var(ϵ^{WIID})	0.00123*** (0.000)	0.00123*** (0.000)	0.217*** (0.020)
Var(ϵ^{WID})	0.000196 (0.000)	0.000196 (0.000)	0.023 (0.014)
Variance Gini	0.0124*** (0.002)	1 (constrained)	1 (constrained)
Cov($\epsilon^{WIID}, \epsilon^{GCIP}$)	0.00378*** (0.000)	0.00378*** (0.000)	0.436*** (0.037)
Cov($\epsilon^{WB}, \epsilon^{GCIP}$)	0.0000770 (0.000)	0.0000770 (0.000)	0.044 (0.032)
Cov($\epsilon^{WB}, \epsilon^{WIID}$)	0.000776*** (0.000)	0.000776*** (0.000)	0.725*** (0.025)
Observations	518	518	518

the total variation of the indicators.

From the estimation, we can derive the unobserved factor or score. In the context of the factor model, there are two standard estimators of the score: Thomson’s regression predictor and Bartlett’s estimator. Bartlett’s estimator was justified as being the minimum variance unbiased estimator of f when Λ and Ψ are known. However, assuming $\Psi > 0$, Neudecker and Satorra (2003) show that Thomson’s predictor has a smaller mean square error than Bartlett’s score. For this reason, we use Thomson’s predictor for the score.

Table 5: Correlation matrix

	Predicted Score	Gini-MIFA	Gini-WIID	Gini-WB	Gini-WID	Gini-GCIP
Predicted Score	1					
Gini-MIFA	0.59	1				
Gini-WIID	0.89	0.45	1			
Gini-WB	0.92	0.48	0.94	1		
Gini-WID	0.90	0.59	0.87	0.90	1	
Gini-GCIP	0.88	0.69	0.77	0.80	0.86	1

Table 5 shows the correlation of the predicted score with each of the inequality indicators. The MIFA indicator has a correlation with the predicted score close to 0.6. The highest correlation is associated with the World Bank indicator. From the analysis in the previous sections and the correlations with other sources of inequality²⁹ we see that this new methodology to calculate inequality using satellite imagery provides a very useful indicator.

5.3 Determinants of the fit of the MIFA index

To complete the characterization of the MIFA index, in this section we analyse the explanatory factors of the difference between the predicted country score and the MIFA Gini. For this purpose, we use variables related to the pixels used for the calculation of the MIFA index (proportion of top-coded pixels, proportions of pixels with no light, etc.) and other characteristics of the country capture by traditional indicators (population, density, GDP per capita, etc.).³⁰

Since we are interested in explaining the size of the difference independently of the sign, we use as a dependent variable the absolute difference between our MIFA Gini and the common factor. Table 6 presents the results. The increase in population density, size of the country or level of development reduces the difference between the predicted score, the common factor Gini, and our MIFA Gini. By contrast, the increase in the percentage of pixels with no light reduces the adjustment. The proportion of pixels that are top-coded does not have a significant effect on the difference, seeming to indicate that our correction for top-coded areas is working, at least when

²⁹Considering also that those indicators are measured with error.

³⁰Lessmann and Seidel (2017) use a similar strategy to predict regional GDP per capita using data on night-light by region.

Table 6: Determinant of the quality of the adjustment: country data

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Ln Population density	-0.04*** (0.006)	-0.05*** (0.007)	-0.04*** (0.006)	-0.04*** (0.006)	-0.04*** (0.006)	-0.04*** (0.005)	-0.04*** (0.005)
Ln Number of pixels in the country (size)		-0.01** (0.005)	-0.01** (0.004)	-0.01** (0.004)	-0.01*** (0.004)	-0.008* (0.004)	-0.008* (0.004)
Ln GDP/capita			-0.05*** (0.004)	-0.02*** (0.008)	-0.03*** (0.008)	-0.03*** (0.008)	-0.03*** (0.008)
1 - Percentage of no light pixels				0.17*** (0.041)	0.15*** (0.042)	0.14*** (0.035)	0.13*** (0.035)
Percentage of top-coded pixels					-0.89*** (0.169)	0.08 (0.143)	0.09 (0.146)
Constant	0.66*** (0.02)	0.88*** (0.08)	1.34*** (0.08)	1.04*** (0.11)	1.09*** (0.11)	0.80*** (0.12)	0.83*** (0.12)
Region Fixed Effect	-	-	-	-	-	Yes	Yes
Year Fixed Effect	-	-	-	-	-	No	Yes
R^2	0.11	0.12	0.39	0.42	0.43	0.61	0.62
Observations	518	518	518	518	518	518	518

the regional dummies are included.

6 Applications

In this section we consider several applications of the data constructed in the previous sections. In particular, we consider the calculation of ethnic and regional inequality. Recently several papers have proposed to measure these dimensions of inequality using methodologies based on nighttime light data.

Alesina et al. (2016) use data on night-light to calculate a Gini index across ethnic groups for each country. Their methodology implies calculating the mean income per capita, proxied by luminosity per capita, at each ethnic homelands, and ranking the average income by ethnic group to calculate the Gini index. This index has the same value independently of the proportion that each group represents in the population of the country. In addition Alesina et al. (2016) also calculate what they define as spatial inequality, which aggregates luminosity per capita across large equally sized, or Thiessen, polygons.

Lessmann and Seidel (2017) propose a method to calculate regional income inequality using nighttime light data. They argue that luminosity cannot be used directly to measure regional income per capita. For this reason they construct a regression model to predict GDP per capita using several variables including the average nighttime light within a region, or the number of top coded pixels and dark pixels in each region as well as fixed effects for satellite configuration. Using a random effects model they predict regional GDP per capita. Once they have the income per capita of each region they calculate several measures of inequality (e.g. Gini, Theil) weighted by population in each region.

As we described before, our procedure to calculate inequality is different. We use nighttime

light per capita of very small pixels to construct an index of inequality. There are more than 2.2 million populated pixels in the average country, and 81 individuals by populated pixel. Our approach, differently from the previous research, allows calculating inequality within and between regions, and ethnic groups, since we have many observations for each region / ethnic group. In order to perform this exercise we calculate the Theil index since it is a decomposable indicator.³¹

The Theil index can be calculated as

$$(13) \quad Theil = \frac{1}{N} \sum_{i=1}^N \frac{y_i}{\bar{y}} \ln \left(\frac{y_i}{\bar{y}} \right)$$

where N is the number of individuals, y_i is the income of the i-th individual and \bar{y} is the average income. In our case y_i is the average light per capita in each pixel.

The decomposition of the Theil index takes the form

$$(14) \quad Theil = \sum_{i=1}^M s_i T_i + \sum_{i=1}^M s_i \ln \frac{\bar{y}_i}{\bar{y}} = Within + Between$$

where M is the number of groups, T is the Theil index for group i, \bar{y}_i is the average NTL per capita for group i, and s_i is the proportion of NTL per capita of group i over the total NTL per capita of the country, $s_i = \frac{N_i \bar{y}_i}{N \bar{y}}$.

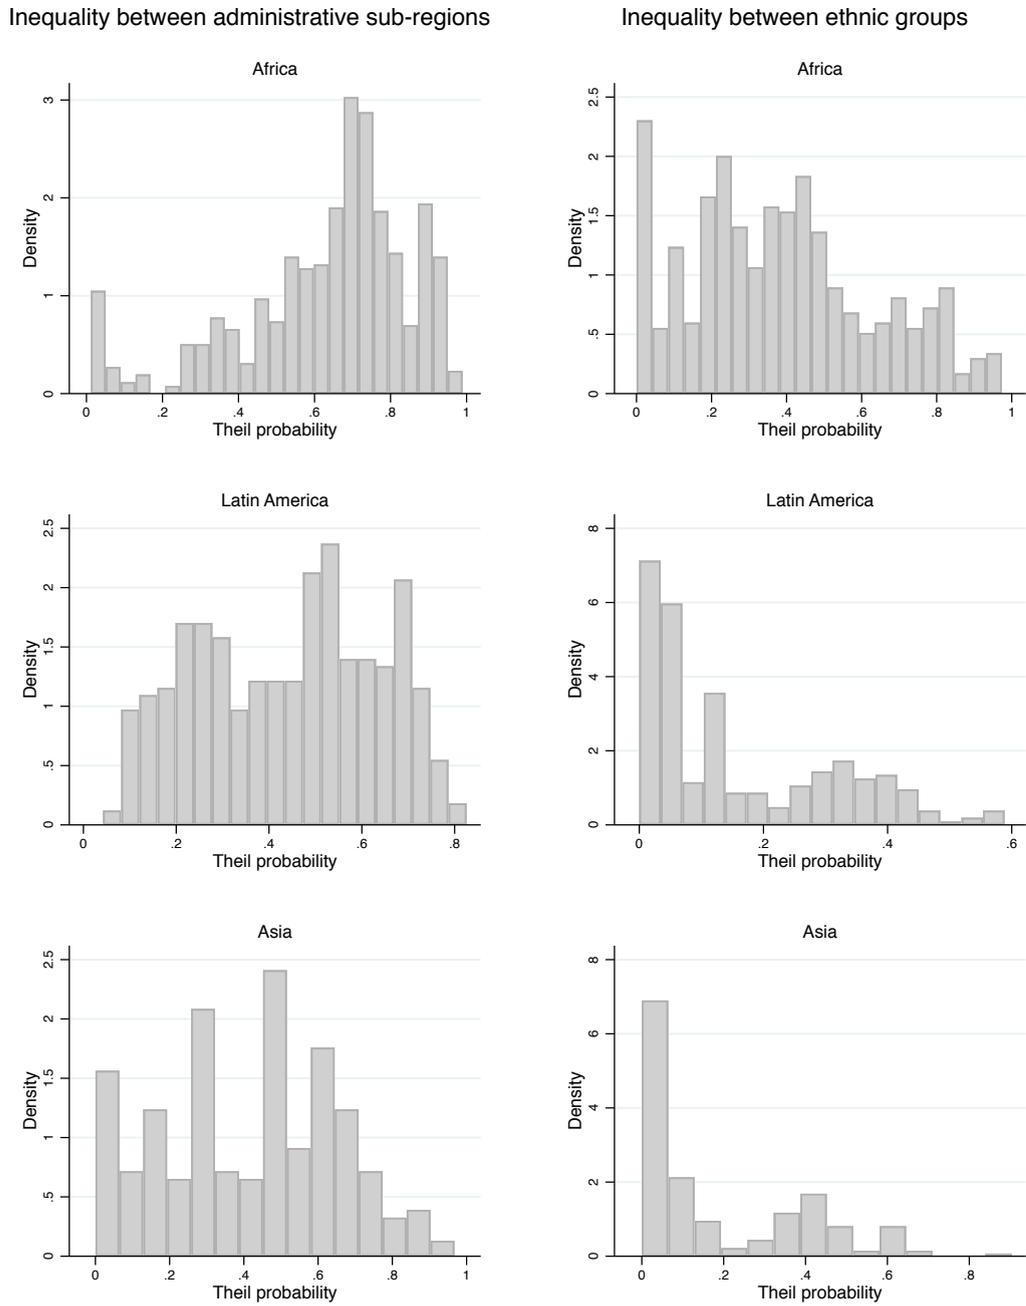
We use the previous methodology to calculate within and between regions/ethnic groups across all the countries in the world. To define the regions for each country we use the political borders provided by The Global Administrative Unit Layers (2010) –GAUL– project of the Food and Agriculture Organization (FAO) of the United Nations. It compiles and standardizes the best available information on different administrative units for all the countries in the world, providing a contribution to the standardization of the spatial dataset representing administrative units. GAUL always maintains global layers with a unified coding system at country, first (e.g. departments) and second administrative levels (e.g. districts). Where data is available, it provides layers on a country by country basis down to third, fourth and lower levels. Because GAUL works at global level, unsettled territories are reported. At the first administrative level there are 3.304 regions.

For the ethnic borders we use the ‘Geo-referencing of ethnic groups’ (GREG) dataset, which relies on maps and data drawn from the Atlas Narodov Mira that stems from a major project of

³¹Opposite to the case of the Theil index, which is decomposable, the Gini index is not. For this reason in the decomposition of the Gini index there is an overlap component which occurs when there are individuals in different groups who have an income difference between them that is of opposite sign to the average income difference between the groups. This effect is called trans-variation between groups. This overlap component makes it difficult the interpretation and the comparison across countries.

charting ethnic groups worldwide, to represent groups territories as polygons. Figure 13 shows the distribution of between inequality across administrative regions and ethnic groups in different continents.

Figure 13: Inequality between regions and ethnic groups: countries by continent



Notes- Insert note here.

Figure 13 represents the distribution of the proportion of between inequality over total inequality. Between ethnic groups inequality represents also an important proportion of total inequality in Africa.³² In fact, the six countries with the highest proportion of between ethnic groups inequality over total inequality are from Africa: Central African Republic (82%), Chad (75%), Kenya (64%), D.R. of Congo (64%), Tanzania (63%) and Gabon (54%).³³ The inequality between regions is also an important source of inequality in Africa. The highest proportion of between regions inequality is reached in Guinea Bissau, Mauritania, Uganda, Burundi, Central Africa Republic and Gambia.³⁴ Between regions inequality in Latin America is quite important, while the inequality between ethnic groups is less relevant. That is also the case in Asia: the inequality between regions is more relevant than the inequality between ethnic groups.

7 Conclusions

We propose a new method to calculate inequality measures using nighttime light (NTL) captured by satellite imagery. The economic literature has shown the usefulness of these measures to proxy economic activity. In this paper, we show that these data can also be useful to measure economic inequality around the world using a common methodology. Currently, there are several indicators of inequality across countries produced by different institutions. They use different data and methodology, producing indicators that do not necessarily coincide. Our objective is to propose a procedure that can be generalised to a global scale and produce methodologically homogeneous measurements that could be applied to any geographical scale. Our methodology is based on constructing small pixels and calculating the average NTL per capita as a proxy for the income per capita of the representative individual of the pixel. This procedure produces an average of 2.2 million populated pixels per country, and an average population of 81 individuals per pixel. Then, we rank the pixels and perform the calculation of the Gini index following the standard formulation. We name this index the MIFA (Measuring Inequality from Above) Gini index. To check the appropriateness of this indicator we run a common factor analysis across the Gini indices available from different sources. The exploratory analysis confirms that all those indices reflect one common factor. The confirmatory analysis provides the loading of each indicator on the common factor, and the predicted score of the common inequality factor. Despite the difficulties in matching country Gini indices constructed using other sources, the MIFA Gini index works quite

³²Alesina et al. (2016) also find that Africa is the most ethnically unequal place in the world.

³³The procedure Alesina et al. (2016) only includes Chad among the countries with highest cross-ethnic-group inequality.

³⁴As in the case of between ethnic groups inequality, the ranking of countries with the highest levels of spatial inequality following the procedure of Alesina et al. (2016) only has one coincidence: Central African Republic.

well producing a correlation with the common inequality factor close to 0.6. The MIFA index works best for countries with a low proportion of zero light pixels, high population density, and high income per capita. Our methodology naturally allows calculating within and between groups inequality for each country. We applied this decomposition to two dimensions: administrative regions and ethnic groups. Future research will further investigate the relevance of this inequality decomposition, in particular the ethnic dimension, to explain conflict, public good provisions, etc.

References

- Abrahams, A., Oram, C., and Lozano-Gracia, N. (2018). Deblurring DMSP nighttime lights: A new method using Gaussian filters and frequencies of illumination. *Remote Sensing of Environment*, 210:242–258.
- Addison, D. M. and Stewart, B. P. (2015). Nighttime lights revisited : the use of nighttime lights data as a proxy for economic variables. Technical report, The World Bank.
- Alesina, A., Michalopoulos, S., and Papaioannou, E. (2016). Ethnic inequality. *Journal of Political Economy*, 124(2):428–488.
- Atkinson, A. B., Piketty, T., and Saez, E. (2011). Top incomes in the long run of history. *Journal of Economic Literature*, 49(1):3–71.
- Bluhm, R. and Krause, M. (2018). Top Lights Bright Cities and their Contribution to Economic Decelopment. Technical Report December, CESifo.
- Chen, X. and Nordhaus, W. D. (2011). Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences*, 108(21):8589–8594.
- Deaton, A. (2005). Measuring poverty in a growing world (or measuring growth in a poor world). *Review of Economic and Statistics*, 87(1):1–20.
- Dobson, J. E., Bright, E. A., Coleman, P. R., Durfee, R. C., and Worley, B. A. (2000). LandScan: A global population database for estimating populations at risk.
- Elvidge, C., Baugh, K., Dietz, J., Bland, T., Sutton, P., and Kroehl, H. (1999). Radiance calibration of DMSP-OLS low-light imaging data of human settlements. *Remote Sensing of Environment*, (68):77–88.

- Elvidge, C., Ziskin, D., Baugh, K., Tuttle, B., Ghosh, T., Pack, D., Erwin, E., and Zhizhin, M. (2009). A Fifteen Year Record of Global Natural Gas Flaring Derived from Satellite Data. *Energies*, 2(3):595–622.
- Elvidge, C. D., Baugh, K. E., Anderson, S. J., Sutton, P. C., and Ghosh, T. (2012). The night light development index (nldi): A spatially explicit measure of human development from satellite data. *Social Geography*, 7:23–35.
- Henderson, J. V., Storeygard, A., and Weil, D. N. (2012a). Measuring economic growth from outer space.
- Henderson, M., Yeh, E. T., Gong, P., Elvidge, C., and Baugh, K. (2003). Validation of urban boundaries derived from global night-time satellite imagery. *International Journal of Remote Sensing*, 24(3):595–610.
- Henderson, V., Storeygard, A., and Weil, D. (2012b). Measuring economic growth from outer space. *American Economic Review*, 102 (2):994–1028.
- Hsu, F. C., Baugh, K. E., Ghosh, T., Zhizhin, M., and Elvidge, C. D. (2015). DMSP-OLS radiance calibrated nighttime lights time series with intercalibration. *Remote Sensing*, 7(2):1855–1876.
- Imhoff, M. L., Lawrence, W. T., Elvidge, C. D., Paul, T., Levine, E., Privalsky, M. V., and Brown, V. (1997). Using nighttime DMSP/OLS images of city lights to estimate the impact of urban land use on soil resources in the United States. *Remote Sensing of Environment*, 59(1):105–117.
- Jean, N., Burke, M., Xie, M., Davis, W. M., Lobell, D. B., and Ermon, S. (2016). Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301):790–794.
- Keys, R. (1981). Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 29(6):1153–1160.
- Lanscan (2018). Methodology of landscan.
- Lessmann, C. and Seidel, A. (2017). Regional inequality, convergence, and its determinants—a view from outer space. *European Economic Review*, 92:110–132.
- Mellander, C., Lobo, J., Stolarick, K., and Matheson, Z. (2015). Night-time light data: A good proxy measure for economic activity? *PLoS ONE*, 10 (10).
- Michalopoulos, S. and Papaioannou, E. (2013). Pre-Colonial Ethnic Institutions and Contemporary African Development. *Econometrica*, 81(1):113–152.

- Michalopoulos, S. and Papaioannou, E. (2014). National institutions and subnational development in Africa. *Quarterly Journal of Economics*, 129(1):151–213.
- Montalvo, J. and Reynal-Querol, M. (2020). Ethnic diversity and growth: revisiting the evidence. *Review of Economics and Statistics*, forthcoming.
- Neudecker, H. and Satorra, A. (2003). On best affine prediction. *Statistical Papers*, 44:257–266.
- Piketty, T. and Saez, E. (2014). Inequality in the long run. *Science (New York, N.Y.)*, 344(6186):838–43.
- Pinkovskiy, M. and Sala-I-Martin, X. (2016). Lights, camera . . . income! Illuminating the national accounts-household surveys debate. *Quarterly Journal of Economics*, 131(2):579–631.
- Pinkovskiy, M. and Sala-i Martin, X. (2016). Lights, camera... income! illuminating the national account-household surveys debate. *Quarterly Journal of Economics*, pages 579–631.
- Pinkovskiy, M. and Sala-i Martin, X. (2020). Shinning a light on purchasing power parities. *American Economic Journal: Macroeconomics*, 12 (4):71–108.
- Rosenbloom, T. (2009). Crossing at a red light: Behaviour of individuals and groups. *Transportation Research Part F: Traffic Psychology and Behaviour*, 12:383–394.
- Rozenfeld, H., Rybski, D., Gabaix, X., and Makse, H. (2011). The area and population of cities: new insights from a different perspective on cities. *American Economic Review*, 101 (5):2205–2225.
- Sen, A. and Foster, J. (2005). *On Economic Inequality*. Clarendon Press.
- Small, C., Elvidge, C. D., Balk, D., and Montgomery, M. (2011). Spatial scaling of stable night lights. *Remote Sensing of Environment*, 115(2):269–280.
- Small, C., Pozzi, F., and Elvidge, C. D. (2005). Spatial analysis of global urban extent from DMSP-OLS night lights. *Remote Sensing of Environment*, 96(3-4):277–291.
- Sutton, P., Roberts, D., Elvidge, C., and Meij, H. (1997). A comparison of nighttime satellite imagery and population density for the continental united states. *Photogrammetric Engineering and Remote Sensing*, 63:1303–1313.
- Wu, J., He, S., Peng, J., Li, W., and Zhong, X. (2013). Intercalibration of DMSP-OLS night-time light data by the invariant region method. *International Journal of Remote Sensing*, 34(20):7356–7368.

Yitzhaki, S. and Schechtman, E. (2013). The Gini Methodology: A Primer on a Statistical Methodology. *Springer Series in Statistics. New York and Heidelberg: Springer, 2013, pp. xvi, 548.*

A Pixel level

A.1 Nighttime light

We use the cloud-free night-time light (NTL) that provides artificial light on the Earth obtained from the Defense Meteorological Satellite Program-Operational Line Scanner (DMSP-OLS). This data offers a long-run viewpoint approach to the human economic activity from 1992 and 2013. The DMSP-OLS values ranges from 0 (non-illuminated areas) to 63 (strongly illuminated areas). DMSP-OLS has been extensively used in economics. It presents some limitation due to its relatively coarse spatial resolution, lack of on-board calibration system, and relatively low radiometric resolution. The combination of these limitation result in signal saturation of OLS nighttime observations (top-coding).

To overcome this problem, Elvidge et al. (1999) proposed the combination of the DMSP-OLS data set with auxiliary data set obtained from the pre-flight sensor calibration. As a result, it is possible to have a radiance calibrated NTL free of top-coding. Despite the case of the OLS radiance calibrated NTL might potentially solves the saturation bias, there are some important caveats. First, as the pre-flights calibration are seldom done, the data is only available in seven years (1996,1999,2000,2003,2004,2006 and 2010). Second, because the DMSP-OLS does not have any on-board calibration device, there is not a clear baseline value from the actual level of saturation. Thus, the radiance values should still be considered relative and not absolute (Hsu et al., 2015). Third, there are important differences in the time frequency with year (fewer orbits) and geographical coverage for the calibration values might varies across time. As a result, the radiance calibrated products present high instability in its quality and coverage (Hsu et al., 2015).

From 2013, a new version of cloud-free NTL data was released by the Earth Observation Group (EOG) in the National Geophysical Data Center (NGDC) of the National Oceanic and Atmospheric Administration (NOAA), using the Suomi National Polar-orbiting Partnership (Suomi-NPP) satellite using the Visible Infrared Imaging Radiometer Suite (VIIRS) sensors. VIIRS-NPP data sets are available at an annual scale for 2015 and 2016, and a monthly scale for the 2012–2019 period. Unlike DMSP-OLS data, the VIIRS data do not have any over-saturation bias and a higher resolution (15 arc-second, approximately 500 m) but it is available for a short period of time. For this reason we do not consider these data.

A.2 Population Raster

The best data available on the human population distributions are the national population census. Nonetheless, this information is rarely available at individual level and its time frequency is non-homogeneous across countries. National census data are only available at census units (usually at villages/cities, districts or, in some cases, sub-district level), which are generally quite aggregated geographical units. As a consequence, there are different techniques that decomposed the census units into smaller geographical units using different allocation models. Thus, in order to create high-resolution population raster it is necessary to re-allocate people within each census units. There are three main alternatives to do this procedure:

Without Allocation Model: This method only relies on the census unit to allocate population into a specific grid. By assuming that people is homogeneously distributed within the census unit, it possible to obtain a pixel-based population approach. When census unit is larger than the resolution, population is allocated based on the proportion geographical area.

Lightly Modeled Allocation Model: Under this method, people is allocated within the census unit using external information about the actual location of people. Normally, the external information is luminosity which provides information of the actual location of people during night. Then, the allocation within the census unit is based in the share of the total luminosity.

Modeled Allocation: This approach uses a complex interpolation model that based on several layers (e.g. roads, coastlines, elevation, among others) predicts the most likely allocation with the census unit. Then based, on this predictions the population raster is built.

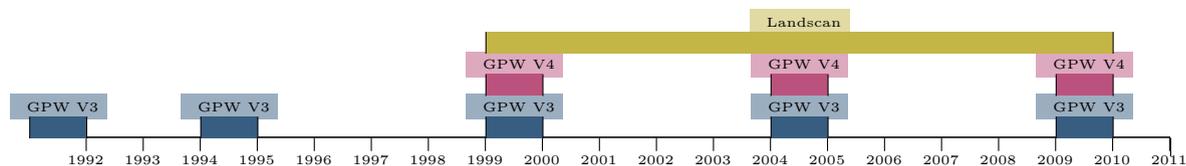
The main sources of information on population raster currently available are described in table 7.

Table 7

Population Raster	Resolution	Census Units	Reallocation Method
GPW V.3	2.5 arc-min (5kmx5km)	399,747	Unmodeled
GPW V.4	30 arc-sec (1kmx1km)		Unmodeled
GRUMP	30 arc-sec (1kmx1km)	12,500,00	Lightly Modeled Allocation Model
Landscan	30 arc-sec (1kmx1km)	8,285,172	Modeled Allocation

Figure 14 presents the temporal availability of the population raster by the different sources.

Figure 14: Temporal availability by population rasters



A.2.1 Landscan

Landscan is often known as the daytime population as it combines a multi-layered, asymmetric, spatial modeling approach (also known as “smart interpolation”), to reallocate population based on the census units. This modeling process is based on the census units for each country and primary geospatial input or ancillary datasets, including land cover, roads, slope, urban areas, village locations, and high resolution imagery analysis (Landsat, 2015). This information is available yearly from 2000-2015.

A.2.2 The Gridded Population of the World

The Gridded Population of the World (GPW) is known as the night-time population as it relies on the information housing without taking into account any other economic variable. Currently there are available two versions with different resolutions: Version 3 (2.5 arc-min) and Version 4 (30 arc-sec 1kmx1km). The sooner is available for 1990,1995,2000,2005 and 2010. Data from 1990 and 2000 are based on round census data for each country. The remaining days are built using projections, which are corrected to guarantee that does not exceed UN projections. In contrast, besides to provide a much higher resolution, the latter uses either 2000 and 2010 census rounds to build gridded information on total population,sex, urban/rural designation.

A.2.3 Other sources

Another known source of high-resolution population data is the Global Rural-Urban Mapping Project –GRUMP– a 30 arc-second (1km at the equator) resolution that include a relocation available for 1990, 1995 and 2000. Despite its high resolution, the allocation of population within the grid is based on luminosity.

B Country data

B.1 Gini indices

B.1.1 PovcalNet/World Development Indicators (WDI)

Host institution: World Bank.

Source data: National household surveys and Luxembourg Income Study (LIS).

Coverage: 1,000+ income or expenditure household surveys from 159 countries with different coverage depending on the year. Comparative estimate for 91 countries from 1993 to 2013, 81 countries for 2008 to 2013 (mostly developing countries) for trend data. In our database, available data from 2000 to 2013.

Description: The indicator is mainly based on household data. In particular, expenditure rather than income surveys for developing countries and income surveys for high-income countries with Gini indexes calculated from the LIS database. No interpolation is done.

Available for download through the World Bank Open Data website: <https://data.worldbank.org>.

B.1.2 World Income Inequality Database (WIID)

Host institution: UNU-Wider.

Source data: Household survey statistics obtained from national statistical offices, the Socio-Economic Database for Latin America and the Caribbean (SEDLAC), the OECD Income Distribution database (IDD), the EU-Statistics on Income and Living Conditions (EU-SILC), LIS and PovCal-Net. Version 3.4 of January 2017.

Coverage: 182 countries for various years starting in the 1940s. Most data series start in the 1980s. In our database: available data from 2000 to 2013.

Description: Collection of country-year estimates from many databases. The inequality statistics are categorized by income concept and equivalence scale. No interpolation is done.

Available for download through the Wider-UNU website: <https://www.wider.unu.edu>.

B.1.3 World Wealth and Income Inequality Database (WID.world)

Host institution: World Inequality Lab.

Source data: Fiscal (income tax) data and data from national accounts combined with other sources (household income and wealth surveys, inheritance and wealth tax data, as well as wealth rankings published in the media). Indicator: pre-transfers, national income

Coverage: National accounts and fiscal data from 57 countries for the years 1980-2016 and older

data for some countries. In our database: available data from 2000 to 2013.

Description: Annual estimates of the distributions of income and wealth (“distributional national accounts”) harmonized based on definitions of income and wealth that are consistent with the macroeconomic national accounts. Pareto interpolation is done.

Available for download through the Wid World website: <https://wid.world>.

B.1.4 Global Consumption and Income Project (GCIP)

Arjun Jayadev, Rahul Lahoti, Sanjay G. Reddy.

Source data: The project builds on various existing data resources, including the LIS.

B.2 GDP per capita

GDP per capita is gross domestic product divided by midyear population. GDP is the sum of gross value added by all resident producers in the economy plus any product taxes and minus any subsidies not included in the value of the products. It is calculated without making deductions for depreciation of fabricated assets or for depletion and degradation of natural resources. Data are in current U.S. dollars. Source: World Bank (<https://data.worldbank.org>).

B.3 Population density

Population density is midyear population divided by land area in square kilometers. Population is based on the de facto definition of population, which counts all residents regardless of legal status or citizenship—except for refugees not permanently settled in the country of asylum, who are generally considered part of the population of their country of origin. Land area is a country’s total area, excluding area under inland water bodies, national claims to continental shelf, and exclusive economic zones. In most cases the definition of inland water bodies includes major rivers and lakes. Source: World Bank (<https://data.worldbank.org>).

B.4 Region

Common geographical classification used by World Bank to assign a region to each country. There exist seven categories: East Asia and Pacific, Europe and Central Asia, Latin America and Caribbean, Middle East and North Africa, North America, South Asia, Sub-Saharan Africa.