Top Lights - Bright Spots and their Contribution to Economic Development

Richard Bluhm (Leibniz University Hannover, Germany) and Melanie Krause (Universitat Hamburg, Germany)

Paper prepared for the 34th IARIW General Conference

Dresden, Germany, August 21-27, 2016

Session 5 (Plenary): New Approaches to Studying the Causes and Consequences of Poverty, Inequality, Polarization, and Social Conflict

Time: Thursday, August 25, 2016 [Morning]

Top Lights

Bright Spots and their Contribution to Economic Development

Richard Bluhm*   Melanie Krause†

VERY PRELIMINARY. PLEASE DO NOT CITE OR DISTRIBUTE.
JULY 31, 2016

Abstract

Satellite data on nighttime luminosity are an increasingly popular proxy for economic activity, but their utility for analyzing comparative development across the globe is severely limited by top-coding. The commonly used satellites do not accurately capture the brightness of large and densely populated cities. As a result, they underestimate differences between urban and rural regions, as well as developed and developing countries. Our main contribution is to propose a new and easy-to-use procedure to correct for top-coding of the lights data. We show that, like top incomes, top lights are Pareto distributed. On this basis, we derive simple formulas for the top-adjusted mean and spatial Gini coefficient. Furthermore, we develop simulation methods to correct the data at the pixel level. Our top-coding correction raises the worldwide Gini coefficient of spatial inequality in lights by about 9 percentage points. This rather large increase underlines the importance of big cities for global economic activity. We present further applications to show that top-coding affects estimates of the income elasticity of light, urban-rural differences, and regional or ethnic inequalities.

JEL Classification: D3, O1, O18, C34

Keywords: Development, inequality, nighttime lights, top-coding, top incomes

*Leibniz University Hannover, Institute of Macroeconomics, Königsworther Platz 1, 30167 Hannover, Germany; Maastricht University; UNU-MERIT; Data-Pop Alliance, E-mail: *bluhm@mak.uni-hannover.de*

†Hamburg University, Department of Economics, Von-Melle-Park 5, 20146 Hamburg, Germany; E-mail: *melanie.krause@wiso.uni-hamburg.de*

1 Introduction

Economic activity in countries with low quality statistical data is hard to track. This problem not only plagues studies focusing on developing countries but also those dealing with growth, inequality and comparative development on a global scale. Advanced economies tend to be over-represented in empirical studies, while less developed countries, especially in Africa, are often left out due to data unavailability. This data constraint is so severe that it has been referred to as a 'statistical tragedy' (Devarajan, 2013). As a result, we know very little about spatial inequality or polarization in developing countries at the sub-national level, although there is a presumption that they matter in theory (Esteban and Ray, 1994; Duclos et al., 2004; Anderson, 2004; Anderson et al., 2012).

Light emissions observed by weather satellites circling the earth at night are increasingly seen as a way out of this statistical dilemma. A growing literature in economics now uses nighttime lights as a proxy for national or local economic activity (e.g. see Chen and Nordhaus, 2011; Henderson et al., 2012; Michalopoulos and Papaioannou, 2013; Hodler and Raschky, 2014). At first sight, the advantages of night lights are obvious. They are publicly available as a yearly panel from 1992 onwards for nearly all parts of the world. They have a high resolution compared to regional accounts data and are measured uniformly across the globe. Hence, they are usually considered to be comparable both within and between countries. Henderson et al. (2012) show that night lights can predict output growth at the country level while allowing us to circumvent thorny discussions over adjustments for exchange rates and price levels. One important drawback of this new data, however, is that they are top-coded in big cities and densely populated areas. Due to sensor saturation, the bustling center of a metropolis appears no brighter than that of a mid-sized town, distorting estimates of regional inequality and convergence.

The main contribution of this paper is to establish that top-coding matters for many important research questions, demonstrate that the upper tail of the distribution of light intensities at night follows a Pareto law, and present top-coding corrected estimates of average lights and spatial inequalities. We outline a simple procedure for simulating the top tail of the lights data and geo-referencing the simulated data at the pixel level. Ultimately, this will allow us to present a top-coding corrected data set of light intensities at the national, regional and pixel level over the period from 1992 to 2013 for most parts of the populated world.

Geo-referenced images of night lights are typically obtained from the National Geophysical Data Center (NGDC) at the National Oceanic Administration Agency (NOAA). Satellites from the Operational Linescan System of the Defense Meteorological Satellite Program (DMSP-OLS) have been orbiting the earth for some decades now with the primary purpose of detecting sunlit clouds. As a byproduct, they measure light emissions in the evening hours between 8:30 and 10:00 pm local time around the globe every day. The recorded data are pre-processed (removing observation of cloudy days and sources of lights which are not man-made, such as auroral lights or forest fires) as well as averaged over cloud-free days. The result is an image of annual light intensities from 1992 to 2013 for every pixel around the globe at a resolution of 30 by 30 arc seconds[1] (approximately 0.86 square kilometers at the equator). Figure 1 shows how the night lights provide a view on economic activity and human settlement patterns around the world just before the turn of the millennium.

Figure 1: Map of 'stable lights' in 1999, saturated

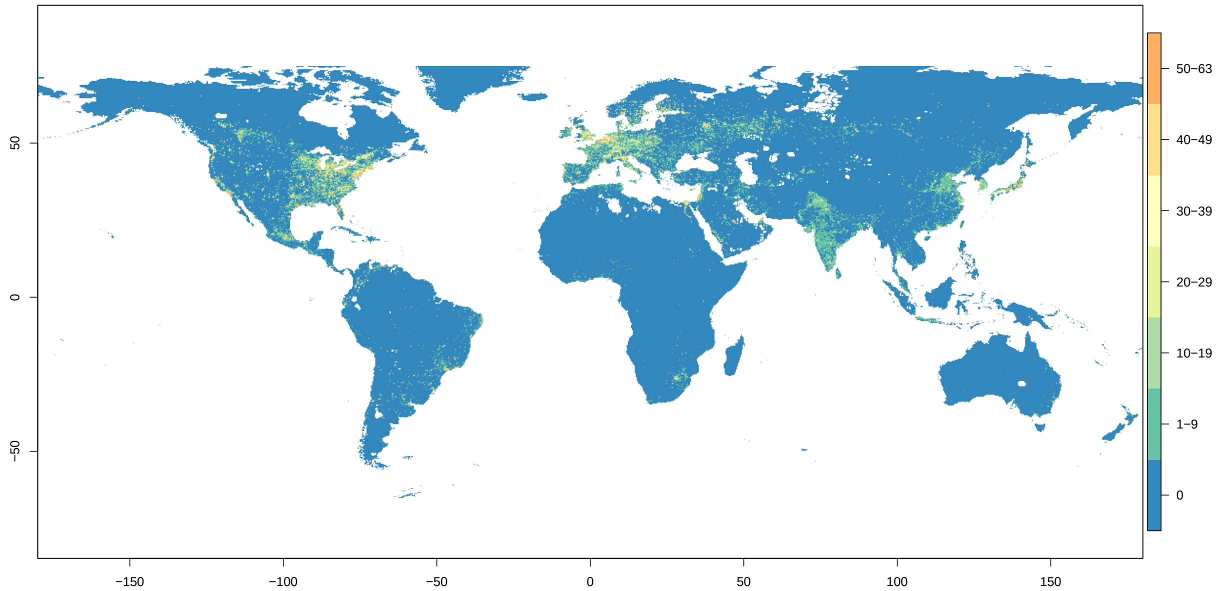

Notes: Illustration of the nighttime lights data. See Elvidge et al. (2009) for details on how the data are constructed. 'Saturated' refers to version 4 of the stable lights DMSP-OLS annual composite.

The light intensity values of this so-called 'stable lights'[2] product are recorded in a fixed range of digital numbers (DN) from 0 (missing or completely dark) to 63 (bright). Sensor saturation implies that the satellites are not able to capture a light intensity higher than 63 DN. As a result, they are unable to distinguish between a mid-sized city and a booming metropolis in most developed countries. This top-coding of the data understates differences in the light intensity of rural and urban areas, induces downward bias in inequality measures, and artificially raises the speed of convergence. As we will show, it is an under-appreciated problem for studies of regional convergence on a global scale (such as Leßmann and Seidel, 2015), light-based estimates of national GDP growth (Henderson et al., 2012), analyses of the economics of urban agglomerations (Storeygard, 2016), and estimates of spatial inequalities (Alesina et al., 2016).

[1]Areas close to the polar zones (65 degrees south and 75 degrees north latitude) are usually excluded. As these regions are very sparsely populated, the exclusion affects approximately 0.0002 percent of the global population (see Henderson et al., 2012).

[2]We use the terms 'stable lights' and 'saturated lights' interchangeably to refer to the same data.

Globally, 5-7% of all lit pixels can be considered to be top-coded and the scale of the truncation is enormous. Cities like New York, Tokyo or Berlin are easily more than ten times brighter than recorded by the stable lights data. As an example, consider Figure 2 for Germany in 1999, where the recorded values hit the 63 DN threshold in most larger urban areas. The saturated data on the left are not able to differentiate among larger cities. The unsaturated lights on the right clearly allow us to locate the brightest spots (e.g. Berlin, Hamburg, Frankfurt, Munich) and distinguish them from dimmer cities. Top-coding tends to affect developed countries more than their less-developed counterparts. However, as bustling economic centers such as Jakarta and Lagos grow further, it will also lead to distorted results within developing countries. In fact, top-coding may resolve part of the puzzle presented by the fact that night lights are more strongly correlated with economic activity in developing rather than developed countries (see Pinkovskiy and Sala-i Martin, 2016; Nordhaus and Chen, 2015). This conjecture was recently supported by Mellander et al. (2015). Using high resolution grids of Swedish administrative data, they find that night lights predict local activity in the form of population, wages and establishments much better once top-coding is taken into account. We will show that this point carries over to measures of spatial inequalities, estimates of the national light elasticity of income in OECD countries, and sub-national light-output elasticities in Germany.

Figure 2: Germany in 1999, saturated and unsaturated lights

(a) Saturated **(b)** Unsaturated

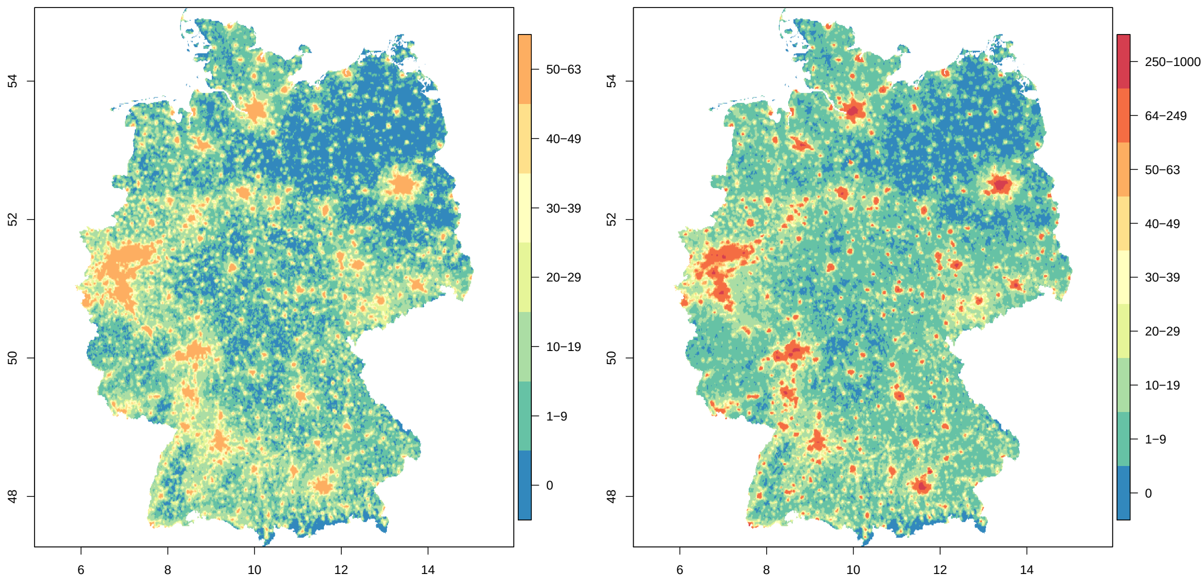

Notes: Illustration of the saturated and unsaturated data for Germany. Panel a) shows a map of Germany based on the stable lights data from satellite F121999 that is commonly used in the literature. Panel b) shows the same map using the radiance-calibrated data in 1999 from Hsu et al. (2015). Both data have been binned as shown in the legend and the color scales were adjusted so as to be comparable.

Top-coding has received little attention in the related literature so far, largely because we lacked reliable time-series data on non-saturated lights.[3] We simply did not know how much of the scale is missing. This situation has now changed. For seven years, additional satellites flew with various fixed gain settings that are less sensitive to light and thus capable of capturing the upper part of the light distribution. The resulting 'radiance-calibrated' data are no longer top-coded but still plagued by at least three major problems:[4] (i) The data are only available for a few years whereas the stable lights series forms a panel from 1992 to 2013. (ii) Even for values that are not strictly at the top-coding boundary, there are large differences in brightness between the two available series. (iii) The radiance-calibrated series exhibits great variability over time and is not strictly comparable across images from different satellites and years.

Measurement errors are also present in the 'stable lights' data and limit their reliability in the time series dimension. The satellites' sensors deteriorate over their lifetime and have to be replaced every couple of years. As a result, the images are not strictly comparable *across* and *within* satellites. In settings such as panel regressions, we may resort to a combination of satellite fixed effects and time fixed effects. In other settings, it is not possible to ensure the comparability of the DNs over time.[5] This holds in particular when changes in the shape of the regional, national or global distribution of lights are to be analyzed. Yet even in panel regressions, estimates of the long-run income elasticity of lights tend to greatly exceed estimates of short-run elasticities. This occurs at least in part because the latter are dominated by noise.

We offer a simple and computationally attractive solution to the top-coding problem. Borrowing methods from the top incomes literature, we propose to extend the distribution of lights using a Pareto tail. For income data, which typically suffers from top-coding, modeling the top share based on a Pareto distribution and then recomputing inequality measures has become the *de facto* standard in the literature (e.g. see Atkinson, 2005; Atkinson et al., 2011; Dell, 2005). We derive simple formulas for the spatial Gini coefficient and average light intensity that analytically combine the two data sources. We show that in order to calculate the top-coding corrected Gini coefficient, it is not necessary to adjust the underlying image or have access to the radiance-calibrated data for that matter. For all economic applications at the sub-national level, however, a top-coding correction at the pixel level is required. We show that this can also be done in a straightforward way: First, we use the radiance-calibrated data to infer the shape

[3]Until recently only one cross-section of unsaturated lights in 2006 was made available by the NGDC.

[4]Hsu et al. (2015) outline a procedure to obtain these 'radiance-calibrated' images by blending the various fixed gain images with the stable lights series.

[5]Appendix C provides a primer on different approaches to solve the satellite inter-comparability problem. Elvidge et al. (2009) propose to scale the various images to a reference area, Sicily, and a particular reference satellite. We show that although this method perfectly scales the images, it is conceptually flawed and removes much of the relevant time-series variation. We then propose alternative approaches that will be incorporate in later versions of this paper.

parameter of the Pareto distribution. We then augment the lower part of the observed distribution with an upper part based on simulated draws from the estimated Pareto distribution. In a nutshell, we distribute part of the observed histogram around the top-coding threshold into the usually unobserved right tail of the data. Working with a worldwide sample of 2% of all pixels, we show that this top-coding correction approach works very well. We are currently extending this procedure to the full sample in order to construct a complete top-coding corrected data set with global coverage over the period from 1992 to 2013.

Applying our top-coding corrected Gini formula confirms the conjecture that top-coding makes a big difference.[6] The estimated worldwide Gini coefficient of lights increases, on average, by 9 percentage points from 0.48 to 0.57 throughout the period from 1992 to 2013. In later years the corrections are even larger. At the country level, we see that countries which are (i) richer, (ii) smaller in size and (iii) highly urbanized tend to have particularly high shares of pixels affected by top-coding. The spatial Gini coefficient of these countries increases by up to 20 percentage points after the top-coding correction. Top-coding also matters in developing countries. In countries such as Egypt, where one bright economic hub (Cairo) coexists with many dimly lit areas, top-coding substantially distorts estimates of regional inequality. Phrased provocatively, much of the previous literature has been estimating development without cities.

We then turn to three important economic applications to study and assess the impact of top-coding. The first application revisits the seminal paper by Henderson et al. (2012) and shows that the radiance-calibrated data works better for OCED countries. We extend this work to the sub-national level and show that with our corrected series estimates of the light-output are considerably more robust and double in size. The second application evaluates the economic significance of urban areas and capital cities as in Storeygard (2016) [tbc]. Preliminary results for Germany confirm that top-coding rises with population density and that our corrected data recovers a sensible economic ranking of Germany cities. Our final application turns to ethnic inequalities as in Alesina et al. (2016) and revisits their link with underdevelopment [tbc].

The paper proceeds as follows. Section 2 illustrates the extent of top-coding around the world. Section 3 establishes that the upper tail of the night lights is in fact Pareto distributed and presents our top-coding correction. In Section 4 we present preliminary results of the impact of top-coding and calibration on important applications. Section 5 concludes. The appendices provides additional tables and proofs, as well as a primer on between and within satellite measurement errors.

[6] We focus on the Gini coefficient as it is the most widely used inequality index. In future versions of this paper we are planning to analyze the impact of the top-coding correction on other inequality indices. Furthermore, it will be insightful to see how measures of polarization are affected where cities are likely be influential as well (Esteban and Ray, 1994; Duclos et al., 2004; Anderson et al., 2012).

2 The Extent of Top-coding Around the World

How important is top-coding anyway? To answer this question, we first have to distinguish between lit and dark regions. In fact, a cursory look from outer space (such as Figure 1) shows that vast parts of the world are completely dark. If we consider a rather coarse grid that is often used in the literature, where each cell is 0.5×0.5 decimal degrees (about 55×55 km at the equator), then 49% of all cell in 2010 do not contain a single lit pixel. Small and densely populated countries (Israel and Belgium) have a much higher share of lit cells than those with (mostly) uninhabited wide plains (USA and Brazil). In any analysis of global growth, distribution or development, these uninhabited areas without a trace of light will be excluded. Hence, we restrict our attention to the lit regions of the earth.[7]

Next, we have to settle on an appropriate unit of analysis. Most applications work with pre-defined regions or grids, but we are interested in recovering the full distribution of lights and these levels hide all heterogeneity within a region or cell. In our case, working at the highest-available resolution is preferable. The native resolution of the data at the pixel level is 30×30 arc seconds. On a global scale, this is a formidable computational exercise. Each image contains more than 700 million pixels, about half of which are illuminated. To reduce the computational burden, we work with spatial random sample of 2% of all lit pixels within all countries that have a landmass larger than 500 km^2.[8] The resulting data set contains more than 4 million pixels per year from 197 countries and territories. The number of pixels range from 4 in Kiribati to 364,910 in the United States. The average country is represented by 22,601 pixels.

Using this data we can conduct a pixel level comparison in 2010, for which both the saturated 'stable lights' and the unsaturated 'radiance-calibrated' data are available.[9] Both images are annual composites of daily pictures, but the 'radiance-calibrated' data can be composites of the same satellite flown with different sensor settings or even composites of different satellites (for details see Hsu et al., 2015). Note that this makes a one-to-one mapping of the two data sources extremely difficult, if not impossible. This is why we primarily focus on differences in the scale and distribution of the data.

Table 1 shows that the mean pixel luminosities of the unsaturated lights are only slightly higher than the saturated ones. The average lit pixel of the global sample has

[7] One should keep this in mind when interpreting measures of spatial inequality: When a hitherto dark spot lights up and becomes part of the sample, inequality between lit regions will increase.

[8] In each of 197 countries and territories, 2% of all lit pixels are sampled. Obviously, in countries with fewer lit pixels, this sample represents a smaller share of the whole area than in others. Also, because the 2% sample is a panel it includes pixels which were unlit in previous year. We ignore these values in our computations for lit pixels, which leads to the discrepancy between the numbers of pixel in the data set and the numbers of lit pixels reported in Table 1. We also cross-check our results with a different 1% sample of lit pixels from all countries.

[9] The saturated data are averaged across the whole year, while the radiance-calibrated data come from satellite F16, which circled the earth from 11 January to 9 December 2010.

Table 1: Summary statistics of global lights in 2010

	World	USA	Brazil	Israel	Nigeria	China	Belgium
	Panel a) Grid level – saturated lights						
Number of lit cells	**36183**	3603	1962	23	280	2782	27
Lit cells as % of total	**51.04**	71.30	64.41	100.00	82.84	67.62	100.00
	Panel b) Pixel level – saturated lights						
Number of lit pixels	**4,115,575**	342,246	48,287	22,371	26,838	132,226	2229
Mean light intensity	**17.69**	18.40	17.93	31.19	15.69	17.74	39.23
(Standard deviation)	**(15.76)**	(16.93)	(15.34)	(21.58)	(14.17)	(15.49)	(16.45)
Max. light intensity	**63.00**	63.00	63.00	63.00	63.00	63.00	63.00
Spatial Gini	**0.4344**	0.4490	0.4139	0.3916	0.4146	0.4258	0.2409
	Panel c) Pixel level – unsaturated lights						
Number of pixels	**4,124,784**	342,889	48,143	22,365	26,748	132,725	2,229
Mean light intensity	**19.57**	23.10	19.70	46.74	13.37	19.83	41.21
(Standard deviation)	**(46.09)**	(52.92)	(41.83)	(84.87)	(19.94)	(45.42)	(45.36)
Max. light intensity	**2379.62**	1674.74	648.63	1099.83	459.58	1590.68	398.90
Spatial Gini	**0.6148**	0.6602	0.5980	0.6613	0.4726	0.6053	0.4731
	Panel d) Pixel level comparison						
Share of pixels >= 55 (saturated)	**6.46**	8.49	6.15	24.71	4.60	5.91	24.50
Unsaturated mean if saturated >= 55	**136.33**	149.67	144.59	137.49	75.83	138.47	93.55

Notes: The table reports summary statistics using three different data sources. Panel a) uses the saturated lights at the grid cell level, where each cell is 0.5×0.5 decimal degrees. Panel b) uses a 2% sample of the saturated lights at the pixel level, where each pixel is 30×30 arc seconds. Panel c) uses the same 2% sample of the unsaturated lights. Panel d) compares both sources at the pixel level.

a saturated luminosity of 17.69 DN and an unsaturated luminosity of 19.57 DN. What is striking is the difference in the two scales. The standard deviation of the unsaturated pixels is three times as high as the one of the saturated pixels (46.09 DN vs. 15.76 DN). The scale of the saturated lights ends at 63 DN, while the unsaturated lights do not have a theoretical upper bound. The brightest pixel on earth is recorded at 2379.62 DN rather than 63 DN. Naturally, this results in a much higher spatial Gini coefficient when the unsaturated rather than saturated lights are considered (0.43 vs. 0.61). While not all of the differences between the saturated and unsaturated lights can be ascribed to top-coding, the comparison does give some insight into its scope. In the last two lines we see that 6.46% of all lit pixels fall between 55 and 63 DN, while their unsaturated counterparts are, on average, more than twice as bright (136 DN).

The remaining columns of Table 1 show the corresponding statistics for selected countries. Clearly, top-coding is an issue which affects virtually all the countries of the world, albeit some more than others. Developing countries like Brazil and Nigeria reach the saturated maximum as well; more than 4% of lit pixels are brighter than or equal to 55 DN. Still, countries which are (i) highly developed, (ii) small and (iii) highly urbanized like Israel and Belgium have many more pixels affected by top-coding. There, more than

20% of illuminated pixels have a saturated recorded brightness larger than or equal to 55 DN. This is also reflected in their high mean light intensity which is larger than 30 DN. In Belgium, this high average light intensity coupled with sensor saturation implies a very low Gini in lights (0.24), which doubles when the full scale of the unsaturated lights is considered.[10]

It is tempting to conclude from cursory comparison that researchers should simply use the radiance-calibrated data instead of the saturated data. Unfortunately, they are no panacea. The unsaturated, radiance-calibrated, satellites are only available for 7 rather than 22 years between 1992 and 2013. More worryingly, they contain a lot of instability, noise and (most probably) measurement error. As a case in point, consider Table 2 and Table 3 which compare the variation over time as indicated by both data sources.

Table 2: Summary statistics of the saturated data over time

Satellite	Mean	% with DN >= 55	Gini	Satellite	Mean	% with DN >= 55	Gini
F101992	14.21	4.30	0.4521	F152002	13.39	3.99	0.4557
F101993	12.21	3.51	0.4724	F152003	10.42	2.94	0.5105
F101994	12.02	3.70	0.4898	F152004	10.40	3.06	0.5196
F121994	15.62	5.02	0.4438	F152005	10.69	3.07	0.5206
F121995	13.32	4.16	0.4690	F152006	10.86	3.23	0.5227
F121996	12.99	3.95	0.4720	F152007	11.08	3.12	0.5165
F121997	13.63	4.36	0.4667	F152008	18.07	6.64	0.4331
F121998	13.99	4.69	0.4606	F162004	12.25	3.83	0.4746
F121999	14.81	5.07	0.4548	F162005	10.80	3.17	0.5003
F141997	11.18	3.51	0.5038	F162006	12.73	3.89	0.4782
F141998	11.03	3.43	0.5049	F162007	13.52	4.50	0.4714
F141999	10.24	3.05	0.5048	F162008	13.35	4.37	0.4784
F142000	11.24	3.33	0.4984	F162009	14.00	4.68	0.4779
F142001	11.66	3.77	0.4988	F182010	17.69	6.46	0.4344
F142002	11.99	3.88	0.4994	F182011	15.34	5.63	0.4657
F142003	12.17	4.22	0.5033	F182012	16.95	6.46	0.4427
F152000	13.41	4.08	0.4525	F182013	16.91	6.57	0.4459
F152001	13.14	3.92	0.4576				

Notes: The table reports summary statistics using a 2% sample of the saturated lights at the pixel level, where each pixel is 30 × 30 arc seconds. The naming convention is as follows. F101993 is the data measured by the satellite named F10 in the year 1993. There are a number of years when two satellites flew concurrently, so that there are 35 satellite-years between 1992 and 2013.

Several points stand out from the time series of the saturated 'stable lights': (i) The mean luminosity of all lit pixels fluctuate between 11 and 16 DN in most years with no discernible trend, but some variation across years and satellites. (ii) Between 3% and 6% of all pixels are at the top of the scale between 55 and 63 DN, with a slight increase in later years. (iii) Spatial inequality in lights, as measured by the Gini coefficient for the saturated data, is relatively stable over the years and not very large, mostly fluctuating between 0.43 and 0.50. We can observe a decrease in spatial inequality towards the end of the sample but its implication is not obvious. More sensitive sensors record higher mean

[10] Interestingly, the brightest Belgian pixel "only" has an unsaturated value of 398.90 DN. China, by contrast, has much fewer bright pixels in relative terms (5.91% with saturated values larger or equal to 55 DN) but the brightest pixel reaches an unsaturated value of 1590.68 DN.

light intensities but lower Ginis and *vice versa*. Degradation of the sensors over time is visible as well; later recordings of any particular satellite tend to also be the brightest.

Table 3: Summary statistics of the saturated data over time

Satellite Obs. period	F12 (96) 16 Mar 96 - 12 Feb 97	F12 (99) 19 Jan 99 - 11 Dec 99	F12 (00) 03 Jan 00 - 29 Dec 00	F14 (02) 30 Dec 02 - 11 Nov 03	F14 (04) 18 Jan 04 - 16 Dec 04	F16 (05) 28 Nov 05 - 24 Dec 06	F16 (10) 11 Jan 10 - 9 Dec 10
Mean	18.80	20.31	22.63	25.32	24.59	21.10	19.57
Std. dev.	(52.26)	(59.05)	(63.65)	(68.39)	(66.55)	(52.06)	(46.09)
Minimum	2.29	0.16	0.15	1.89	0.59	1.70	1.78
Maximum	2068.99	4054.14	3367.00	4219.60	3573.55	2450.00	2379.62
% >= 180 DN	1.48	1.95	2.16	2.57	2.44	1.85	1.63
% >= 300 DN	0.69	0.89	1.04	1.22	1.10	0.72	0.59
Spatial Gini	0.6431	0.6599	0.6359	0.6593	0.6608	0.6319	0.6148
Non-zero pixels	3,678,974	4,452,313	4,060,898	3,839,345	4,076,924	3,852,258	4,124,784

Notes: The table reports summary statistics using a 2% sample of the unsaturated lights at the pixel level, where each pixel is 30 × 30 arc seconds. The naming convention is as follows. F12 is the satellite number and the number in parentheses are the last two digits of the year where most data where observed (here 1999). The full observation period is listed below the satellite name.

While the saturated data seem to exhibit measurement errors across satellites and years, the fluctuations of the unsaturated data are much larger (see Table 3). In particular, the maximum luminosity jumps from 2068.99 to 4054.14 within three years and decreases again by a similar amount. This would be hard to explain in economic terms.[11] The Gini coefficients vary between 0.61 and 0.65 and are typically about 15 percentage points higher than those of the saturated lights. Once again, we do observe that the Gini coefficient decreases towards the end of the sample period, this time with out a corresponding rise in mean luminosity. Even though the range of unsaturated values is very large, only around 1% and 2% of all pixels have values higher than 300 and 180 DN. In other words, values in the 1000's do occur, but they are rare.

The saturated and radiance-calibrated lights report very different light intensities for the same pixel. To illustrate this point, we regress the saturated lights on the radiance-calibrated data in all years where both data sources are available.[12] As Table 4 shows, instead of near equivalence, we only find a regression coefficient of around 0.5. The average pixel has a radiance-calibrated luminosity of more than twice the saturated value. Note that we restrict the data range to all nonzero pixels with a luminosity smaller than 55 in order to exclude all cases where the saturated data do not observe a pixel's full brightness. We motivate this choice below.

We draw three conclusions from these comparisons. First, the saturated lights severely

[11]This variability can mostly be attributed to measurement errors introduced when the different fixed gain images and the stable lights images are merged at the NGDC. There are possibilities to inter-calibrate the satellites (Hsu et al., 2015), but they come with strings attached. We discuss this issue in detail in the appendix of this paper.

[12]The radiance-calibrated satellites have slightly different observation periods than the annual saturated lights. For our comparative analysis, we work with the calendar year in which the majority of radiance-calibrated data falls, e.g. year 2006 for Satellite F16 (06) (28 Nov 2005 - 24 Dec 2006).

Table 4: Regression of saturated on unsaturated data

Year	1996	1999	2000	2003	2004	2006	2010
Unsaturated	0.6057	0.5167	0.4514	0.4546	0.3602	0.4945	0.5835
	(0.0002)	(0.0002)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0003)
Constant	4.2693	5.6831	3.4031	3.3627	3.1298	2.8590	8.0547
	(0.0034)	(0.0040)	(0.0031)	(0.0031)	(0.0027)	(0.0029)	(0.0053)
Adjusted-R^2	0.7574	0.6932	0.7443	0.7834	0.7831	0.7881	0.5062

Notes: The table reports OLS estimates of a regression of all pixels smaller than 55 DN of the saturated data on their counterpart in the unsaturated data. The data are a 2% random sample of lights at the pixel level, where each pixel is 30×30 arc seconds.

underestimate the extent of spatial inequalities. Second, while the radiance-calibrated lights lead to spatial Ginis that are around 15 percentage points higher, not all of the difference between saturated and radiance-calibrated lights can be due to top-coding. Third, working with the raw data from the radiance-calibrated satellites is impractical because they appear infrequently and are unstable. We can circumvent their instability by relying only on the shape of their distributions. As we will see in the next section, this leads to a top-coding adjustment which does account for a sizable proportion of the difference between two data sources and produces relatively stable results over time.

Figure 3: Two Histograms of Saturated Lights in 1999

(a) Saturated when DN > 9 **(b)** Saturated when Unsaturated DN > 160

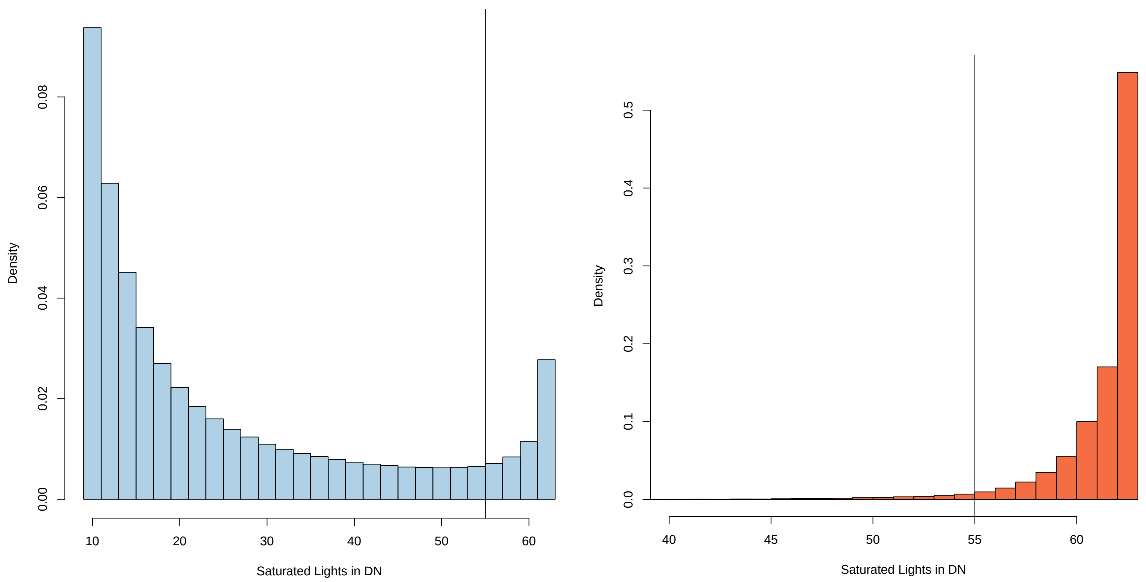

Notes: Illustration of the location of the top coding threshold in the saturated data. Panel a) shows a histogram of the F12 satellite in 1999 for all pixels with a DN greater 9. Panel b) shows a histogram of the same satellite only for pixels where the unsaturated light intensity is greater 160 DN. The input data is a 2% representative sample of all non-zero lights in the stable lights and radiance-calibrated data at the pixel level (see Elvidge et al., 2009; Hsu et al., 2015).

We still have to establish where to put the top-coding threshold. This is more intricate than it looks at first sight. While the scale of the saturated values goes up to 63 DN, we have good reason to assume that many pixels of 62 DN, 61 DN and down to the mid-50s are already subject to top-coding and should be brighter than they are recorded. The reasoning behind this is straightforward. The data providers at NOAA process many daily images into one annual composite. As a result, there may be a considerable amount of pixels below a saturated value of 63 in the annual composite which suffer from top-coding in at least one of the underlying daily images. In fact, Hsu et al. (2015) suggest that this subtle type of top-coding may even start at a brightness as low as 35 DN.

Consider the histogram of the saturated global lights in Figure 3a in 2010. For most of the domain, the histogram is decreasing but this trend reverses in the mid-50s. Many pixels are clustered at saturated light values of about 55 to 63. If only pixels with light intensities of exactly 63 DN were affected by top-coding, then we would expect a declining number of pixels up to 62 DN and a noticeable spike at 63. Clearly, this is not the case. Further evidence along these lines is provided by Figure 3b. The figure shows a histogram of the light intensity of all those saturated lights associated with radiance-calibrated values above 160 (3 times 55 DN). There is a large number pixels with DNs in the 62s, 61s, down to the mid-50s which seem to be top-coded. Hence, we consider a DN of 55 to be the *approximate* start of the range where top-coding plays a significant role (even though a marginal number of pixels below 55 DN may be affected as well). Note that the procedure we suggest will work with any top-coding threshold.

3 Correcting for Top-Coding

3.1 Are top lights Pareto distributed?

Our correction procedure relies on the conjecture that top lights are Pareto distributed. A power law in the tail distribution of nighttime lights has not yet been investigated, but it is an often-found empirical feature in many other fields, ranging from the frequency of words in literary texts over the size of moon craters to, particularly relevant for us, the population of cities (for an excellent overview see Newman, 2005).

A first motivation for a Pareto tail in the distribution of lights comes from the urban economics literature. The upper tail of the size distribution of big cities follows a Pareto law (Gabaix, 1999; Rozenfeld et al., 2014). Such power laws arise when the growth of cities is independent of their absolute size. This is known as Gibrat's Law, if applied to the entire distribution, or Zipf's Law, if attention is restricted to the right tail. Small et al. (2012) use the night lights data to measure the urban extent of cities and empirically demonstrate that there is a Pareto size distribution of large agglomerations. Obviously, top lights come from top cities, so there is a strong analogy with our case. The difference is that we look at individual pixels rather than the cities they form.

A second motivation to look for a Pareto distribution in top lights comes from the well-known fact that top incomes have a power law tail, see for instance Piketty (2003), Atkinson (2005), Atkinson, Piketty, and Saez (2011) and Dell (2005) for the top of the income distribution in individual countries and Lakner and Milanovic (2015) for the global income distribution. While appealing, this analogy has its limits. The physical brightness of a single pixel may not increase as much as the incomes of its residents.

Pareto (1897) first discovered that top incomes above a threshold y_c tend to follow a particular cumulative distribution (CDF)

$$F(y) = 1 - \left(\frac{y_c}{y}\right)^{\alpha} \quad \text{for} \quad y \geq y_c \tag{1}$$

with the associated probability density function (PDF)

$$f(y) = \alpha \cdot y_c^{\alpha} \cdot y^{-\alpha-1} \quad \text{for} \quad y \geq y_c \tag{2}$$

where the parameter $\alpha > 0$. The complement of the CDF in eq. (1) is the survival function; it gives the probability that the random variable Y is larger than the given value y.

$$\overline{F(y)} = \left(\frac{y_c}{y}\right)^{\alpha} \quad \text{for} \quad y \geq y_c. \tag{3}$$

The term 'power law' arises since this may be written as $Cy^{-\alpha}$, where C is a normalization constant. Taking logs delivers the characteristic linear equation

$$\log(\overline{F(y)}) = \alpha \cdot (\log(y_c) - \log(y)) = -\alpha \cdot \log(y) + const. \tag{4}$$

This is the basis of the popular *Zipf plot*, which works as follows: in a log-log-diagram, plot the data on the x-axis and the survival function on the y-axis. If a downward sloping linear relationship emerges, then the data follows a Pareto or power law.

Another unique feature of the Pareto distribution is 'Van der Wijk's Law', which states that the average income (or light) above a given income y is proportional to y. The factor of proportionality is equal to $\frac{\alpha}{\alpha-1} > 1$. More formally, we have

$$\frac{\int_y^{\infty} t f(t) dt}{\int_y^{\infty} f(t) dt} = \frac{\alpha}{\alpha - 1} \cdot y. \tag{5}$$

Figure 4a shows the Zipf plots of the top lights of seven radiance-calibrated satellites for the world sample. We only consider the top of the distribution, starting at 160. Despite a handful of outliers at the very end and some initial curvature, the Zipf plots for top lights look fairly linear and are indicative of a Pareto distribution. Zipf plots often deviate from a line at the very top since fewer and fewer values are observed at the extremes. Sometimes this is somewhat arbitrarily addressed by removing the very top. Instead, we use logarithmic bins so that the size of the bin increase by a multiplicative factor (Newman, 2005). For the Van der Wijk's plot in Figure 4b we also find the expected shape. For each top luminosity value on the x-axis, we plot the average luminosity of all pixels brighter than its value on the y-axis. The relationship looks remarkably linear and the estimated slope is at or slightly larger than 1 for most years. This indicates that 'Van der Wijk's Law' with slope $\alpha/(\alpha-1) > 1$ holds in our data.

Figure 4: Zipf Plot and Van der Wijk's Plot

(a) Zipf Plot **(b)** Van der Wijk's Plot

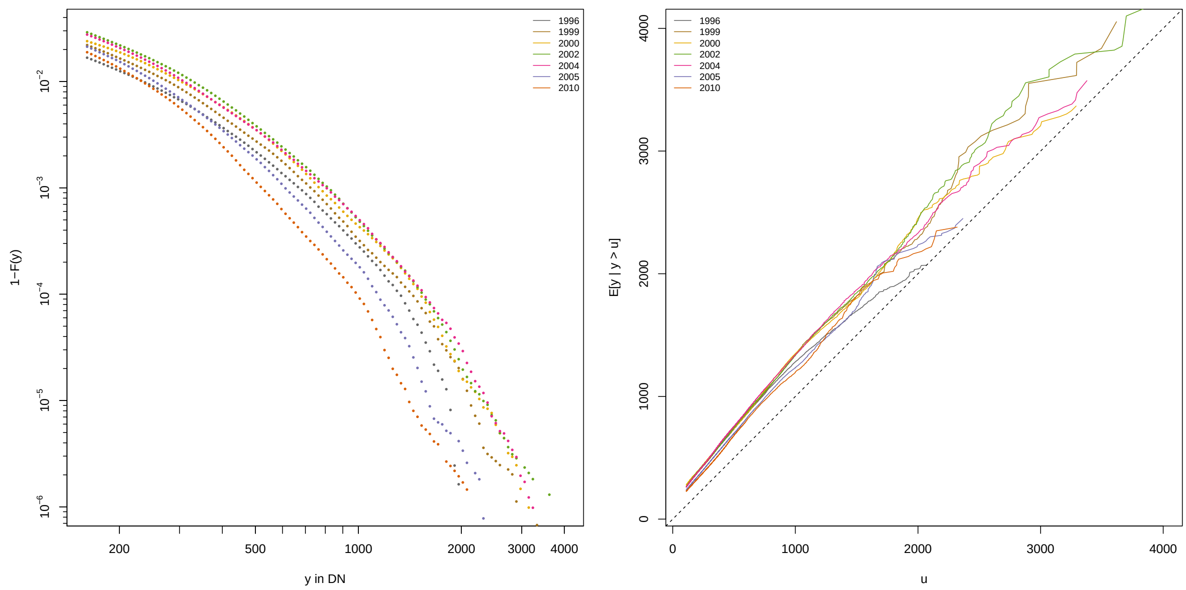

Notes: Illustration of the approximate power law behavior of top lights. Panel a) shows the Zipf plot for all pixel brighter than 160 DN. The figure uses logarithmic binning to reduce noise and sampling errors in the right tail of the distribution (see Newman, 2005). There are about 100 bins in the tail, where the exact number depends on the range of the input data. Panel b) demonstrates Van der Wijk's law, which states that the average light above some value u is proportional to u, this is $\mathrm{E}[y|y > u] \propto u$. Here too the data start at 160 DN. Using a slightly higher or lower threshold value leads to very similar Zipf plots and Mean Excess plots. The input data is a 2% representative sample of all non-zero lights in the radiance-calibrated data at the pixel obtained from Hsu et al. (2015).

Although Zipf plots and related qualitative approaches are standard methods for determining the power law features of observed distributions, there has been a long-standing debate about their appropriateness (see already Lorenz, 1905). Cirillo (2013) demonstrates that data generated by other distributions, such as the lognormal, can look very similar to Pareto distributed data in a Zipf plot. As a remedy, she proposes an additional test which does not rely on the visual evaluation of linearity. The so-called discriminant moment ratio plot shows the coordinate pair of the coefficient of variation (i.e., standard deviation divided by the mean) on the x-axis and skewness on the y-axis. Each parametric distribution has its particular curve of coordinate pairs. We can, for example, divide the area into a Pareto area, a lognormal area and a gray area in between.

Figure 5: Discriminant Moment Ratio Plots

(a) Pixels > 250 **(b)** Pixels > 160

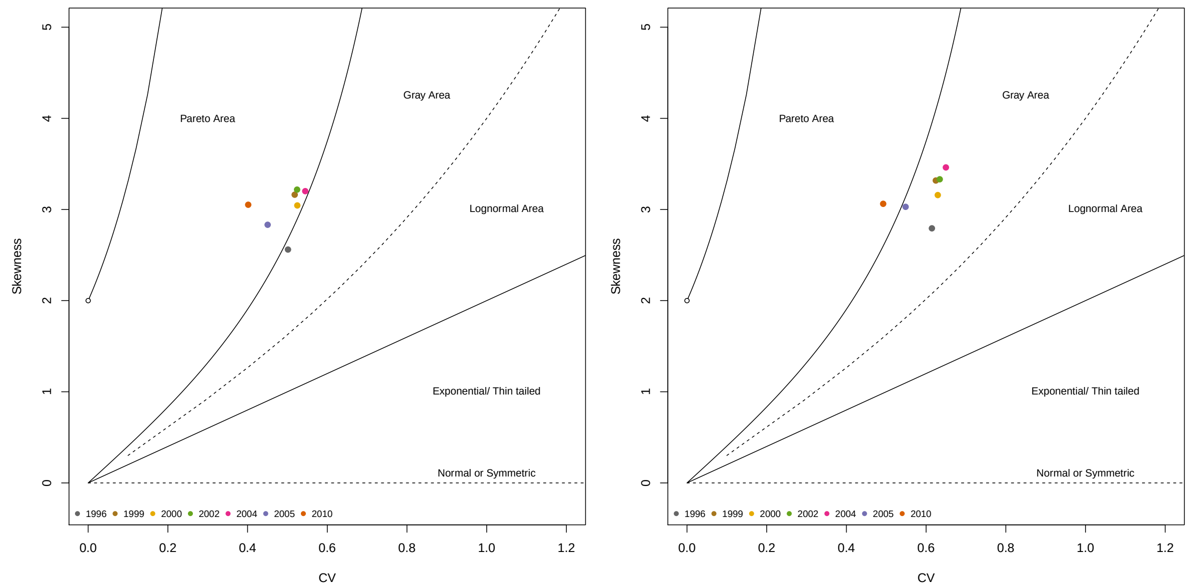

Notes: Illustration of the approximate power law behavior of top lights. Both panels show discriminant moment ratio plots (Cirillo, 2013). Panel a) uses all pixels greater than 250 DN to demonstrate that higher thresholds unambiguously imply a Pareto distribution. Panel b) uses all values greater 160 DN to show that a Pareto distribution approximates this data better than a lognormal distribution. The input data is a 2% representative sample of all non-zero lights in the radiance-calibrated data at the pixel obtained from Hsu et al. (2015).

Figure 5 is our strongest piece of evidence. Figure 5a shows that, for the data with DNs > 250, five out of seven radiance-calibrated satellites are well inside the Pareto area. The other two border the Pareto area but are located in the gray zone close to the Pareto boundary. More importantly, it establishes that the only popular alternative (the lognormal distribution) is not a better representation of our data. This also holds for other thresholds. When we lower the threshold from $y_c = 250$ to $y_c = 160$ (Figure 5b), most of the satellites are not in the Pareto area any more but still firmly in the gray area, far away from the lognormal boundary. It is well-established from the literature on top incomes or the size of cities that the Pareto distribution is often only a good representation of the very top of the relevant distribution, with decreasing fit as the threshold decreases. Overall, our results suggest that there is a strong case to be made for a Pareto distribution of top lights until we go as low as around $y_c = 100$. Our analysis also rules out many other distributions that we omitted from the plot since they are very far from our observations.

3.2 A top-coding corrected mean and Gini coefficient

We propose to correct for top-coding by augmenting the saturated lights with a Pareto tail. All of the saturated lights below the top-coding threshold are kept, while those at or above the threshold are replaced by a Pareto-distributed counterpart. We always use a rule-of-thumb shape parameter of $\alpha = 1.75$ for the Pareto tails. Fixing the parameter may be considered arbitrary but has considerable advantages. First, as panel b) of Table 5 shows, it corresponds to the average shape parameter when we set the Pareto threshold to 160 DN. Second, working with a fixed average Pareto alpha avoids the instability problems of the radiance-calibrated satellites.[13] While it would be interesting to have time series variation in the spatial inequality of top lights only, the scaling and calibration issues of the radiance-calibrated data make identifying an economic trend nearly impossible. A fixed parameter also makes the properties of the new joint data very transparent. While we correct the scale of the data to capture bright cities, changes in inequality will be primarily driven by the vast majority of pixels below the top-coding threshold or by the fact that some places cross the threshold. Nevertheless, from which threshold onwards the Pareto distribution should be fitted is of course debatable. Table 5 offers some alternatives. We later conduct robustness checks with different parameters.

Table 5: Estimation shape parameters with different Pareto thresholds

Satellite	F12 (96)	F12 (99)	F12 (00)	F14 (03)	F14 (04)	F16 (06)	F16 (10)	Average
				Panel a) Threshold $y_c = 140$				
Pareto α	1.4747	1.5691	1.5943	1.5667	1.6024	1.7846	1.8940	**1.6408**
(S.E.)	(0.0044)	(0.0058)	(0.0046)	(0.0043)	(0.0044)	(0.0056)	(0.0061)	**(0.0050)**
				Panel b) Threshold $y_c = 160$				
Pareto α	1.5831	1.6802	1.6957	1.6578	1.7019	1.9188	2.0611	**1.7569**
(S.E.)	(0.0050)	(0.0067)	(0.0054)	(0.0049)	(0.0050)	(0.0066)	(0.0073)	**(0.0058)**
				Panel c) Threshold $y_c = 180$				
Pareto α	1.7004	1.7784	1.8001	1.7456	1.8005	2.0462	2.2175	**1.8698**
(S.E.)	0.0057	0.0076	0.0061	0.0056	0.0057	0.0077	0.0086	**(0.0067)**

Notes: The table reports MLE estimates of the Pareto parameter. The data are a 2% random sample of unsaturated lights at the pixel level above the different thresholds described in panels a) to c). Each pixel is 30×30 arc seconds.

Augmenting the saturated data with a Pareto tail can be done computationally but is not always necessary. We can solve for many interesting features of the full distribution analytically. For example, the top-coding corrected mean luminosity μ of a country or

[13] Rather than working with the single rule-of-thumb Pareto alpha of 1.75, it is of course possible to take seven satellite-specific alpha estimates and transfer them to the corresponding years of the saturated data, using parameter interpolation for intermediate years. This approach gives very similar results overall, albeit with some more jumps in the data, as the noise from the individual radiance-calibrated satellites translates into the Pareto tail. It is also more cumbersome to implement for practitioners.

region is simply the weighted average of the bottom and top means μ_B and μ_T. The latter is the mean of a Pareto distribution starting at the top-coding threshold y_c. Hence, we may write

$$\mu = \omega_B \mu_B + \omega_T \mu_T = \omega_B \mu_B + (1 - \omega_B) \frac{\alpha}{\alpha - 1} y_c \tag{6}$$

where ω_B and $\omega_T = 1 - \omega_B$ are the pixel shares for the data below and above the threshold. In other words, to obtain the top-coding corrected mean luminosity, we can simply plug in the values of ω_B, μ_B, α and y_c without having to replace a single pixel at the top.[14]

When lights are used as a proxy for GDP in studies of inequality and development, another important characteristic of their distribution is the Gini coefficient. Here too, we can find a simple top-coding corrected formula. Given a population with two non-overlapping subgroups (lights below and above the top-coding threshold), the overall Gini can be written as the weighted sum of the bottom-share and top-share Ginis (within-Gini) as well as the difference between the top share of total lights minus top share of pixels (between-Gini). The derivation is relegated to Appendix A but boils down to

$$G = \omega_B \cdot \phi_B \cdot G_B + \omega_T \cdot \phi_T \cdot G_T + [\phi_T - \omega_T], \tag{7}$$

where the shares of all light accruing to the top and bottom groups are $\phi_B = \omega_B \mu_B / \mu$ and $\phi_T = \omega_T \mu_T / \mu$. Plugging the overall mean from eq. (6) as well as the Gini coefficient of a Pareto distribution for the top, $G_T = \frac{1}{2\alpha - 1}$, into eq. (7) and simplifying yields

$$G = \frac{(1 - \omega_T)^2 \cdot \mu_B \cdot G_B + \omega_T^2 \cdot \frac{\alpha}{\alpha - 1} y_c \cdot \frac{1}{2\alpha - 1} + \omega_T \cdot \frac{\alpha}{\alpha - 1} \cdot y_c}{(1 - \omega_T) \cdot \mu_B + \omega_T \cdot \frac{\alpha}{\alpha - 1} y_c} - \omega_T. \tag{8}$$

Now it is straightforward to see that our formula for the top-coding corrected Gini coefficient depends only on the Pareto α (e.g. the rule-of-thumb value 1.75), the top-coded share ω_T (the percentage of pixels above 55 DN in our application), the threshold y_c (55 DN) as well as μ_B and G_B (mean and Gini of the bottom share). All of these values are available without replacement at the pixel level. In fact, with $\alpha = 1.75$ and $y_c = 55$, eq. (8) reduces to[15]

$$G = \frac{(1 - \omega_T)^2 \cdot \mu_B \cdot G_B + 0.4 \cdot 128.33 \cdot \omega_T^2 + 128.33 \cdot \omega_T}{(1 - \omega_T) \cdot \mu_B + 128.33 \cdot \omega_T} - \omega_T. \tag{9}$$

The formula also highlights important qualitative features of spatial inequalities.

[14]A simple numerical illustration shows how correcting for top-coding drives up the mean luminosity. If top-coding starts at $y_c = 55$ DN, affects 5% of the study area of interest, $\alpha = 1.75$ and mean luminosity in the non-top-coded pixels is $\mu_B = 10$ DN, then the corrected mean luminosity is 15.92 DN rather than 12.25 DN.

[15]Note that the rule of thumb alpha of 1.75 implies that the mean of the top share is $\mu_T = \frac{\alpha}{\alpha - 1} \cdot 55 = 128.33$ and the Gini of the top share is $G_T = \frac{\alpha}{2 \cdot \alpha - 1} = 0.40$.

Table 6: Estimates of the global top-coding corrected Gini coefficient

Satellite	Gini unadj.	Top Share ω_T	Mean μ_B	Gini G_B	Gini adj.	Diff.
F101992	0.4521	0.0430	12.14	0.3994	0.5436	0.0915
F101993	0.4724	0.0351	10.46	0.4189	0.5571	0.0847
F101994	0.4898	0.0370	10.16	0.4347	0.5764	0.0865
F121994	0.4438	0.0502	13.25	0.3913	0.5409	0.0972
F121995	0.4690	0.0416	11.27	0.4139	0.5602	0.0911
F121996	0.4720	0.0395	11.05	0.4182	0.5606	0.0887
F121997	0.4667	0.0436	11.50	0.4109	0.5599	0.0931
F121998	0.4606	0.0469	11.70	0.4007	0.5582	0.0976
F121999	0.4548	0.0507	12.37	0.3962	0.5550	0.1002
F141997	0.5038	0.0351	9.39	0.4466	0.5893	0.0855
F141998	0.5049	0.0343	9.28	0.4479	0.5896	0.0847
F141999	0.5048	0.0305	8.66	0.4469	0.5869	0.0821
F142000	0.4984	0.0333	9.55	0.4444	0.5810	0.0827
F142001	0.4988	0.0377	9.75	0.4405	0.5873	0.0884
F142002	0.4994	0.0388	10.04	0.4433	0.5878	0.0884
F142003	0.5033	0.0422	10.04	0.4434	0.5958	0.0925
F152000	0.4525	0.0408	11.42	0.3952	0.5445	0.0920
F152001	0.4576	0.0392	11.22	0.4021	0.5474	0.0898
F152002	0.4557	0.0399	11.44	0.4011	0.5458	0.0901
F152003	0.5105	0.0294	8.92	0.4593	0.5883	0.0778
F152004	0.5196	0.0306	8.83	0.4676	0.5986	0.0790
F152005	0.5206	0.0307	9.12	0.4720	0.5977	0.0772
F152006	0.5227	0.0323	9.21	0.4731	0.6018	0.0791
F152007	0.5165	0.0312	9.52	0.4711	0.5935	0.0770
F152008	0.4331	0.0664	15.06	0.3815	0.5427	0.1096
F162004	0.4746	0.0383	10.33	0.4152	0.5650	0.0904
F162005	0.5003	0.0317	9.18	0.4457	0.5821	0.0818
F162006	0.4782	0.0389	10.81	0.4246	0.5662	0.0880
F162007	0.4714	0.0450	11.32	0.4135	0.5669	0.0956
F162008	0.4784	0.0437	11.20	0.4228	0.5714	0.0930
F162009	0.4779	0.0468	11.73	0.4241	0.5723	0.0944
F182010	0.4344	0.0646	14.74	0.3814	0.5434	0.1090
F182011	0.4657	0.0563	12.65	0.4083	0.5693	0.1036
F182012	0.4427	0.0646	13.95	0.3840	0.5536	0.1108
F182013	0.4459	0.0657	13.84	0.3861	0.5577	0.1118
Average	**0.4786**	**0.0422**	**11.00**	**0.4236**	**0.5697**	**0.0910**

Notes: The table reports the results from the top-coding correction for all satellites at the global level. We report the unadjusted Gini coefficient first, then provide statistics of the top share of pixels, the bottom mean and Gini coefficient, and the new composite Gini coefficient after applying 8 as presented in the text. The last column demonstrates how sizable the correction can be by computing the difference in the two Gini coefficients. All estimates use $\alpha = 1.75$ and $y_c = 55$ for the Pareto-imputation.

A greater share of top-coding, brighter top-coded pixels, and a greater spread in the distribution of the top-coded data all increase the correction in the estimates of inequality. Obviously, these features are more relevant in urbanized or highly developed areas rather than rural areas or less-developed regions.

Table 7: Sensitivity of the global Gini coefficient to the Pareto parameter

Assumed α	1.6	1.65	1.7	**1.75**	1.8	1.85	1.9
Top Gini (G_T)	0.4545	0.4348	0.4167	**0.4000**	0.3846	0.3704	0.3571
Top Mean (μ_T)	146.67	139.61	133.57	**128.33**	123.75	119.71	116.11
Average correction (ΔG)	0.1104	0.1032	0.0967	**0.0910**	0.0859	0.0812	0.0770

Notes: The table reports sensitivity results from the top-coding correction for all satellites at the global level using a range of Pareto coefficients. We only report the difference between the unadjusted and adjusted Gini coefficients, where the new composite Gini coefficients have been computed by applying eq. (8) as presented in the text.

Table 6 applies eq. (9) to the global sample with all the 35 satellite-years from 1992 to 2013. The first column shows the Gini coefficient before the top-coding adjustment (Gini unadj.). The next columns contain the ingredients of the correction (the top share ω_T, the bottom mean μ_B and bottom Gini G_B) which together with $\alpha = 1.75$ result in the corrected Gini coefficient (Gini adj.). Finally, the last column computes the difference in the two Gini coefficients. We see that the global Gini coefficient in lights is on average about 9 Gini points higher after the adjustment, shifting estimates of global inequality in the mid and upper 50s. The top-coding adjustment varies between 7 and 10 Gini points, with larger corrections in later years. Clearly, correcting for top-coding makes a substantial difference. Recall that the top-coding correction only affects a small fraction of lights and the remaining lights are unaltered. *Accounting for the full brightness of only 4% of lights increases the global spatial Gini by 9 percentage points and highlights the contribution of large cities to the distribution of global economic activity.*

Table 7 provides a sensitivity test and considers a range of α parameters from 1.6 to 1.9. Obviously, smaller (larger) α parameters are associated with more (less) inequality in the top and a higher (lower) mean value in the top, leading to a slightly more (less) sizeable top coding correction. But for all these parameters, the average Gini correction over the satellite years remains close to the benchmark correction of 9.1 percentage points (ranging from 7.7 to 11.0 percentage points). Figure 6 plots the time series of the top-coding corrected Gini coefficient in lights from 1992 to 2013, averaging the indices over year where more than one satellite are available. Global inequality in lights increased slightly over the 1990s, remained relatively stable in the first decade of the new millennium and then reached a temporary trough in the aftermath of the global financial crises and great recession (2008-2010). Inequality in lights has since increased again to a Gini coefficient

Figure 6: Time-series of the global top-coding corrected Gini coefficient

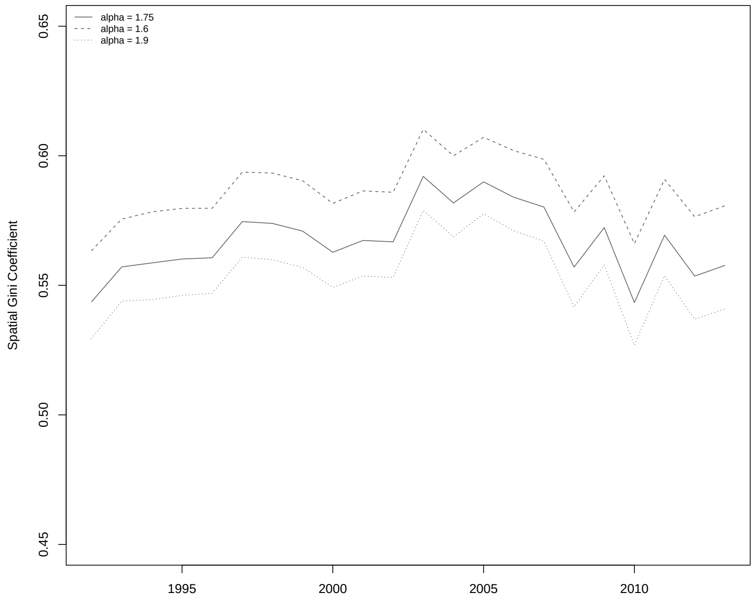

Notes: Illustration of the global top-coding correction. The figure shows global spatial inequality estimated using the formulas presented in the text.

of 0.57. In other words, the level of global spatial inequality is comparable to income inequality in Latin America.

Table 8 shows the the size of our correction for the same countries that we used to motivate the significance of top-coding in 2010. Light inequality in the U.S. is actually representative of the global sample. The top-coding correction increases it from 0.4490 to 0.5744. By contrast, the top-coding correction in Israel and Belgium is much larger. In these two small countries where many pixels are affected by top-coding, our correction increases the Gini coefficient by around 20 percentage points.

Table 8: Top-coding correction for the spatial Gini in 2010

Satellite	Gini unadj.	Top share ω_T	Mean μ_B	Gini G_B	Gini adj.	Diff.
USA	0.4490	0.0849	14.46	0.3834	0.5744	0.1255
Brazil	0.4139	0.0615	15.13	0.3592	0.5212	0.1073
Israel	0.3916	0.2471	21.40	0.3928	0.5811	0.1896
Nigeria	0.4146	0.0460	13.54	0.3612	0.5099	0.0954
China	0.4258	0.0591	15.11	0.3791	0.5296	0.1037
Belgium	0.2409	0.2450	32.69	0.2374	0.4490	0.2081

Notes: The table reports the results from the top-coding correction in selected countries. We report the unadjusted Gini coefficient first, then provide statistics of the top share of pixels, the bottom mean and Gini coefficient, and the new composite Gini coefficient after applying eq. (8) as presented in the text. The last column demonstrates how sizable the correction can be by computing the difference in the two Gini coefficients. All estimates use $\alpha = 1.75$ and $y_c = 55$ for the Pareto-imputation.

What about the magnitude of the top-coding correction in other countries? Table A-1 in Appendix B repeats this analysis for all 197 countries and territories of the globe in 2010. There are very few countries where there is no top-coding correction at all and they tend to either be remote island states (such as Fiji) or extremely poor countries without big cities (such as Guinea-Bissau). At the other extreme, the country with the largest top-coding correction is (unsurprisingly) Singapore. Being a bright and dense city state, 48% of its pixels are subject to top-coding. Taking into account their full brightness drives up Singapore's spatial Gini coefficient by 38 percentage points. Obviously, countries which are poorer, more geographically spread out or less urbanized have fewer pixels affected by top-coding. As a result, the top-coding adjustment increases the Gini coefficient in Japan more than in India (14 vs 6 percentage points). Interestingly, we also observe rather large Gini corrections in some African countries, where urban-rural differences are large (such as a Egypt or South Africa). Bright economic centers such as Cairo and Johannesburg stand out much more in contrast to other parts of the country which are very dimly lit. This highlights an important insight. Although top coding is correlated with the level of development, it is more generally a sign of the dense concentration of economic activity.

3.3 Pixel-level replacement

There many are important questions in development economics and urban economics for which national means and Gini coefficients are not sufficient. For instance, spatial growth regressions at the sub-national level require the corresponding geographical location. More generally, any analysis at the sub-national or grid cell level requires fully geo-referenced data and researchers often want to work with per capita quantities. We now outline how to correct for top-coding in the underlying 'stable lights' image.

Here too, we are going to replace all pixels with values at or above the 55 DN threshold with values sampled from a distribution. We now work with the underlying micro data and do not only correct the summary statistics, so we are concerned with obtaining plausible values for each pixel. It turns out that a truncated distribution is more appropriate in this context. Sampling several thousand lights from a Pareto distribution with $y_c = 55$ and $\alpha = 1.75$ will, by definition, produce a handful of extreme values exceeding 3000 or even 5000 DN. As we discussed earlier, a power law is less evident at the very top, precisely because we rarely observe such extreme lights in the data. There is a natural limit to emitted man-made light in a given place that is not due to gas flares or other non-economic activity. Many of the radiance-calibrated values above 2000 come from Algerian and Saudi-Arabia, pointing to gas flares rather than city centers as their origin. A conservative solution is to sample from a truncated Pareto distribution instead.[16] The truncated Pareto has the CDF

$$F(y) = \frac{1 - \left(\frac{y_c}{y}\right)^{\alpha}}{1 - \left(\frac{y_c}{H}\right)^{\alpha}} \quad \text{for} \quad y_c \leq y \leq H \tag{10}$$

where an upper bound of $H = 1500$ does not affect the overwhelming majority of saturated pixels ≥ 55 to be replaced, but does ensure realistic values at the very top. The brightest pixels in the U.S. and China, for example, have radiance-calibrated values of this magnitude.

The challenge is to geo-reference the probabilistic draws so that the brightest pixels actually end up in the centers of dense urban agglomerations. At first sight, this may be a complicated issue to solve but we can employ a simple trick. Since we know the location of all pixels in both the saturated data and the unsaturated data, we can distribute the simulated values using their combined rank in both data sources. Our suggestion is to first rank all pixels above 55 DN according to the saturated data and then re-rank all pixels with the same value according to their value on the radiance-calibrated image. Usually, this produces a near unique ranking that successfully allocates bright draws to bright locations. We essentially transfer the spatial autocorrelation structure from the upper parts of the radiance-calibrated image to the new composite image, while making sure that only pixels above the top-coding boundary are replaced. The auto-correlation structure is updated each time a new radiance-calibrated image becomes available.

Putting this all together, we propose the following procedure:

1. For each of the 35 satellite-years t of saturated data, calculate the number X_t of pixels ≥ 55 DN.

2. Produce a ranking of these X_t pixels. Rank the pixels based on their saturated values (55-63). For pixels with the same saturated value, rely on the radiance-calibrated data associated with the given satellite or the data from the closest year.

3. Take X_t random draws from a truncated Pareto distribution with the rule of thumb $\alpha = 1.75$, the top-coding threshold $y_c = 55$ and upper bound $H = 1500$.[17]

4. Repeat this procedure 1000 times and average across these 1000 replications. These are the X_t data to be used for the replacement.

5. Replace the X_t saturated pixels ≥ 55 so that the saturated pixel with the i-th highest rank from (2) is replaced by the i-th highest sample value.

[16] The truncated Pareto distribution is also used widely in insurance modeling of natural disasters, Pisarenko and Rodkin (see 2010).

[17] In fact, in order to ensure that tied pixels in the ranking will be replaced by the same sampled value and still use all the sampled values, we sample X_t minus the number of ties in the ranking.

We are still working on solving the computational problems involved in applying this correction to the whole world (each of the 35 available satellite images has more than 700 million pixels). However, preliminary results using a complete time series of all pixels in Germany suggest that this procedure works very well.

Figure 7: Germany, Pixel level replacement, Satellite F12 in 1999

(a) Histogram: Saturated v. Replaced **(b)** Map: Top-Coding Corrected Data

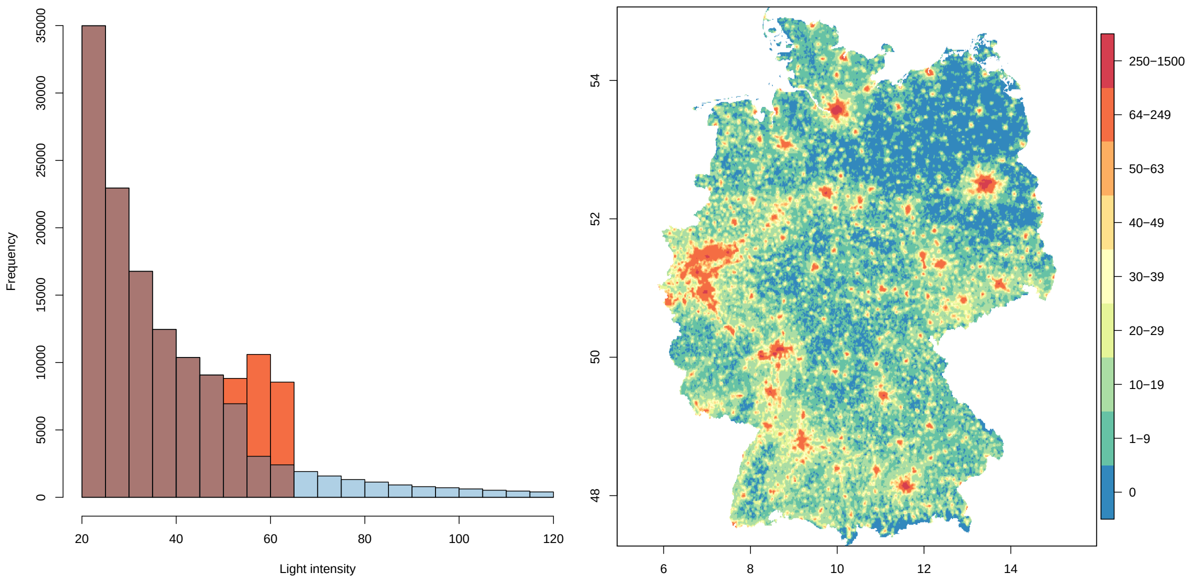

Notes: Illustration of the pixel level correction. Panel a) shows a histogram of all pixels brighter than 19 DN before and after the Pareto interpolation of the top-coded pixels. Panel b) shows a map of Germany after the Pareto interpolation has been applied at the pixel level. The results are computed by sampling from a truncated Pareto distribution as described in the text.

Consider Figure 7a using the satellite F12 in 1999. While the histogram of the saturated data ends with the typical cluster of values between 55 and 63, the replaced values reach far beyond and the histogram declines smoothly. Using the corrected data we can now clearly discriminate among different bright spots as Figure 7b shows. The centers of Berlin, Hamburg, Frankfurt and other German metropolitan areas are distinctly brighter than their surroundings. They are also considerably brighter than second-tier cities. In addition, while the Ruhrgebiet – Germany's largest urban agglomeration with a combined population of more than 10 million – as a whole is rather bright, one can still make out brighter economic centers within, such as Düsseldorf and Essen. In the applications below we explore value added of these corrected images in the cross-section and for panel data. Repeating this pixel level replacement for the whole world at once (or each country at a time) will result in a time-series of top-coding corrected images.

4 Applications

We illustrate the economic significance of top-coding and by examining three prominent research questions. The first application revisits the seminal paper by Henderson et al. (2012) which established that night lights are a good proxy for GDP growth at the national level. We extend this work by also examining this relationship at the sub-national level in Germany. The second application explores a central question in the urban economics literature, namely can we quantify urban-rural differences using nighttime lights? Storeygard (2016), for example, assesses the influence of transport costs on urban growth in Africa using the lights data. Finally, we we study regional and ethnic inequality inspired by a recent path-breaking paper by Alesina et al. (2016).

The aim of these varied applications is to uncover for which questions the top-coding problem is most severe and whether the average or the spread of the distribution are the key area of concern. Our intention is not to second-guess the findings of the original studies but instead to revisit them in this new context.

4.1 Lights and GDP growth

National level estimates: We begin by reproducing the results in Henderson et al. (2012). Specifically, we build a matched-sample of the stable lights data and the radiance-calibrated data for the seven years they have in common over the period from 1996 to 2010. Henderson et al. (2012) calculate average light intensity in each cell that falls on land, weighted by the size of that particular cell, and then run fixed effects regressions of log GDP at constant local prices from the World Development Indicators on their measure of log lights per square kilometer. Weighing each cell by its land area is necessary since the actual area represented by each 30 by 30 arc seconds cell varies due to the curvature of Earth. Henderson et al. (2012) report an income elasticity of lights that fluctuates around 0.26 to 0.28. When we use the 1992 to 2008 sample of the stable lights series, then we also obtain an estimate of 0.282.

Does radiance-calibration change the income elasticity of lights? Table 9 suggests that there seems to be quite some variability in the output-lights relationship. The estimated elasticity already falls substantially by examining a different time period and using less data. Radiance-calibration then induces another drop by about four points. However, these estimates are still well within two standard errors of the results in Henderson et al. (2012). Table 9 also reports per capita elasticities and shows that there is little substantive change in the relationship at the country-level if we are interested in average living standards instead. Here too, the coefficients obtained by using the radiance-calibrated data fall by a similar amount when we use per capita values.

The key point of using the non-saturated data is that it should be better able to capture the growth experiences of rich countries which are relatively more affected by

Table 9: Income elasticity of lights, 1996-2010, country-level

	Saturated Data		*Radiance-calibrated Data*	
	(1) GDP	(2) GDP per capita	(3) GDP	(4) GDP per capita
Lights per km^2	0.226***		0.181***	
	(0.072)		(0.059)	
Lights per capita		0.223***		0.179***
		(0.062)		(0.053)
Constant	25.777***	13.940***	25.749***	13.245***
	(0.014)	(0.945)	(0.022)	(0.785)
Within-R^2	0.708	0.518	0.699	0.503
Observations	1353	1353	1353	1353
Countries	198	198	198	198

Notes: The table reports panel FE estimates. Lights per capita and lights per km^2 are measured in logs. All columns include country and time fixed-effects. Country-clustered standard errors in parentheses. Significant at: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

top-coding than poorer countries. Since we have not yet applied our correction approach to the entire world at the pixel level, we investigate this question by contrasting estimates of the income elasticity of lights for OCED and non-OCED countries obtained using the two different data sources.

Table 10 illustrates an interesting finding. It builds a simple statistical test of whether the OECD and non-OECD elasticities are the same by interacting the lights data with an OECD dummy. Column (1) shows that we reject the hypothesis that the relationship is the same in OECD countries using the stable lights data. In fact, the elasticity for OECD countries is only 0.0045 and a cluster-robust test does not reject the null of zero (with p-value of 0.93). Using per capita value results in a slightly larger elasticity for OECD countries but here too the result is not different from zero (with a p-value of 0.147). Interestingly, the estimate for non-OCED countries only rises a little in return.

The picture is completely different using the radiance-calibrated data. Now we can no longer reject the hypothesis that the elasticities are different in OECD and non-OECD countries. Column (3) of Table 10 shows that there is little change in the elasticity outside of OCED countries as indicated by a statistically insignificant interaction effect. Column (4) then repeats this exercise with the per capita data, where there is even less evidence of a difference between OECD and non-OECD countries. Together these estimates seem to indicate two important insights. On the one hand, the radiance-calibrated data is indeed better suited to analyze growth in richer regions. On the other hand, once the somewhat lower elasticity is taken into account, the light-output relationship at the national level does not seem to differ systematically between rich and poor countries.

Table 10: Income elasticity of lights, OECD v non-OECD, 1996-2010, country-level

	Saturated Data		*Radiance-calibrated Data*	
	(1) GDP	(2) GDP per capita	(3) GDP	(4) GDP per capita
Lights per km^2	0.242***		0.184***	
	(0.072)		(0.062)	
OECD $\times$ Lights per km^2	-0.237***		-0.044	
	(0.035)		(0.058)	
Lights per capita		0.239***		0.180***
		(0.064)		(0.056)
OECD $\times$ Lights per capita		-0.172***		-0.003
		(0.042)		(0.055)
Constant	25.843***	13.769***	25.763***	13.241***
	(0.012)	(0.907)	(0.014)	(0.734)
Within-R^2	0.718	0.526	0.699	0.503
Observations	1353	1353	1353	1353
Countries	198	198	198	198

Notes: The table reports panel FE estimates. Lights per capita and lights per km^2 are measured in logs. All columns include country and time fixed-effects. Country-clustered standard errors in parentheses. Significant at: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Sub-national level estimates: We now turn to sub-national level estimates of the light-output relationship and a first assessment of our corrected series. Here we use data from Leßmann et al. (2015), who compiled GDP per capita estimates and other statistics at the district level in Germany over the period from 2000 to 2011. To this we add our estimates of lights per capita from the saturated and from Pareto-imputed data based on the pixel level correction for Germany. We now take special care to understand whether we add information that is useful in the cross section, in the panel dimension, or both.

The extent to which lights adequately capture local welfare and output is still being actively debated. Mellander et al. (2015), Weidmann and Schutte (2016) and several others show that lights correspond reasonably well with local survey-based measures of welfare in both developed and developing economics. Bickenbach et al. (2016), however, contend that the light-output elasticity at the sub-national level is very unstable in Brazil, India but also the United States and Western Europe. Germany is a good case in point. Leßmann et al. (2015) show that the light output elasticity using the saturated data is substantially lower at the sub-national than at the national level and even becomes indistinguishable from zero when only a few relevant controls are added (such as population and area). Estimates of GDP per capita are unlikely to be the culprit, since the quality of Germany's local data is generally considered to be high.

Table 11: Income elasticity of lights, Germany, NUTS-3 level

	Dependent Variable: GDP per capita					
	Saturated Data			*Pareto-imputed Data*		
	(1)	(2)	(3)	(4)	(5)	(6)
Lights per capita	0.195***	0.136	0.158*	0.210***	0.282***	0.355***
	(0.019)	(0.093)	(0.090)	(0.018)	(0.098)	(0.097)
Population		0.219***	0.266***		0.252***	0.314***
		(0.041)	(0.041)		(0.039)	(0.041)
Area in km^2		-0.126**	-0.122**		-0.022	0.016
		(0.061)	(0.059)		(0.066)	(0.064)
Urban			-0.159***			-0.166***
			(0.043)			(0.042)
Old FRG			0.241***			0.249***
			(0.021)			(0.022)
Adjusted-R^2	0.217	0.435	0.526	0.260	0.445	0.542
Regions	412	412	412	412	412	412

Notes: The table reports cross-sectional OLS estimates. All columns include a constant (not shown). Lights per capita, population and area are all measured in logs. Urban and old FRG are binary variables. The data have been averaged over the period from 2000 to 2011. Robust standard errors in parentheses. Significant at: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11 tackles this issue by examining a cross-section of averages for all 412 districts in Germany over the period from 2000 to 2011. Columns (1) to (3) use the saturated data that is typically used in the literature, columns (4) to (6) mirror the results with our Pareto-imputed data. It is immediately apparent that the informational value of the saturated lights data drops as population, area, an urban dummy and an old federal republic dummy[18] are added to the regression. In fact, columns (2) and (3) suggest that there is little value added in using night lights to predict local output once population and the spatial extent of the region are accounted for. In column (2) the coefficient of lights per capita is no longer significant at conventional levels while population and area are highly significant. This is unfortunate but the question if local lights add much information on top of local population has been echoed elsewhere (Cogneau and Dupraz, 2014). Now consider columns (4) to (6) with our top-coding corrected data. In column (4) the coefficient is comparable to the one obtained using the saturated data (it is within two standard errors of the original estimate). In columns (5) and (6) the sub-national light output elasticity rises substantially and remains highly significant even in the presents of additional controls. Moreover, the area of the sub-national unit is no longer significant,

[18]The old FRG dummy is included, since the former GDR has benefited substantially from public investments in unified Germany and also adopted energy-saving lights differently than the former West.

while the coefficient on population remains significant. In a nutshell, our Pareto-imputed data without top-coding correlate much better with local output. Instead of falling, the cross-sectional light-output elasticity reaches levels which are comparable to its national counterpart from Henderson et al. (2012) – a finding that is completely new to the literature.

Does this result carry over to changes in light and output over time? Table 12 shows a comparable set of results using the complete data over the period from 2000 to 2011. Columns (1) to (3) use the saturated data, while columns (4) to (6) mirror each specification using our Pareto-imputed data. The pattern is once again striking although there are interesting subtleties. Column (1) includes state and time fixed effects. Clearly, the light-output elasticity is once again only about a third of its national counterpart and only significant at the 10%-level. The situation improves only a little once an urban dummy at the district level is included. Column (4) to (5) show that the Pareto-imputed data recovers the same results as in the cross-section when we use only within state variation.

Table 12: Income elasticity of lights, Germany, NUTS-3 level

	Dependent Variable: GDP per capita					
	Saturated Data			*Pareto-imputed Data*		
	(1)	(2)	(3)	(4)	(5)	(6)
Lights per capita	0.103*	0.119**	0.044***	0.220***	0.245***	0.046***
	(0.057)	(0.058)	(0.007)	(0.061)	(0.062)	(0.007)
Population	0.220***	0.271***	-0.554***	0.241***	0.295***	-0.557***
	(0.038)	(0.039)	(0.077)	(0.035)	(0.038)	(0.077)
Area in km^2	-0.142***	-0.154***		-0.058	-0.065	
	(0.038)	(0.039)		(0.041)	(0.041)	
Urban		-0.132***			-0.140***	
		(0.043)			(0.043)	
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	No	Yes	Yes	No
Municipality FE	No	No	Yes	No	No	Yes
Within-R^2	0.599	0.611	0.794	0.606	0.619	0.794
Observations	4944	4944	4944	4944	4944	4944
Regions	412	412	412	412	412	412

Notes: The table reports panel FE estimates. Lights per capita, population and area are all measured in logs. Urban is a binary variable. NUTS-3-clustered standard errors in parentheses. Significant at: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

An interesting result emerges once we allow for municipality level fixed-effects. Now the estimates in columns (3) and (6) using the two different data sources are almost identical. The light output elasticity drops to slightly below 0.5% but remains highly significant. This exemplifies the heterogeneity of the local light-output relationship found in Bickenbach et al. (2016) and Leßmann et al. (2015), since the within-district panel estimates can viewed as stacked time-series regressions with variable intercepts. Our top-coding correction makes little difference in this case because it more strongly influences the district ordering in each cross-section than changes in the light intensity of any particular district over time. Contrary to Bickenbach et al. (2016), we do not think that this necessarily is bad news for researchers interested in using night lights at the local level. Measurement errors, such as those discussed in Appendix C, increase once the within variation of comparatively small units is being used. The annual time variation in average light intensities at the district level both data sources is only 0.2 DN and not representative of the overall scale of the underlying data. Hence, while it may become empirically difficult to robustly trace out the light-output relationship at local level using within variation only, it does not mean that lights do not correlate well with local output and welfare as we have demonstrated here (or others have demonstrated elsewhere Mellander et al., 2015; Weidmann and Schutte, 2016). In any case, our top-coding correction either substantially improves the light-output relationship or at the minimum provides comparable answers.

4.2 Urban-rural differences

We expect that that top-coding will be a major issue for estimating urban-rural differences, precisely because here the cross-sectional comparison matters. Using the same data for Germany from the previous subsection, we can already establish two fundamental facts. First, as expected, top-coding rises with increasing population density (urbanization). Second, the economic ranking of cities using the saturated data is counter-intuitive but a sensible ranking is restored after our correction.

Figure 8 plots the average light intensity per square kilometer over the period from 2000 to 2011 from the saturated and corrected series over the average population density all districts in Germany. Areas officially classified as rural have an average population density of about 300 people per km^2, whereas urban areas have a density of about 1250 people per km^2. Clearly, the two series begin to diverge after a density of about a thousand people per km^2. The saturated series display an exponential decay towards an average intensity of 100 DN[19] whereas the corrected series is approximately *linear* in population density. Another notable feature is that the four brightest cities according to the stable lights data are all medium-sized cities with populations (well) below half a million (Herne,

[19]Note that the saturated series in lights per square kilometer is not bounded by 63 DN because most pixels in Germany are smaller than 1 km^2 but bigger than 500 m^2. Hence the theoretical upper bound is slightly above 100 DN, eg. 63 DN/0.5 km^2 = 126 DN/km^2. This occurs because the plate carrée projection keeps the area of each pixel constant in degrees not kilometers.

Figure 8: Light intensity versus population density in German regions

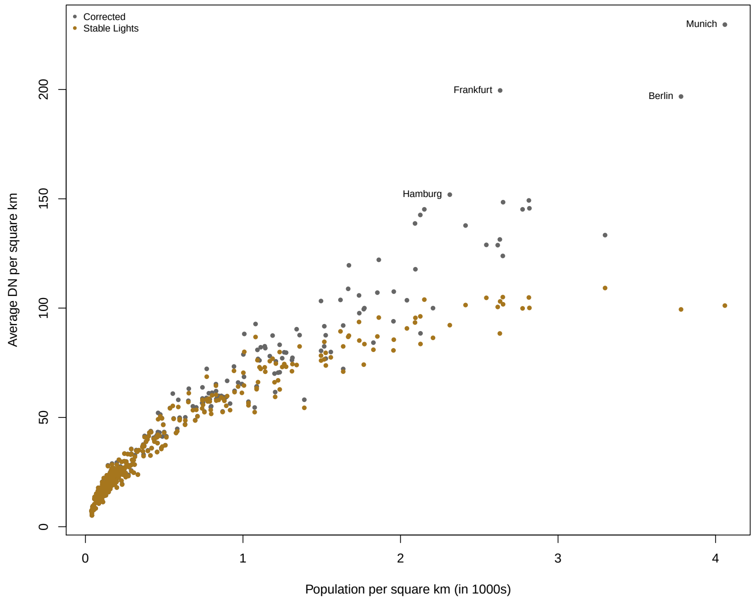

Notes: Illustration of the value added of the top-coding corrected data for urban economics. The data are cross-sectional averages of lights per km^2 and population density in German NUTS-3 regions over the period from 2000 to 2011.

Bochum, Oberhausen, Gelsenkirchen) whereas the corrected data perfectly identifies the four largest and most populated economic centers as the four brightest (Munich, Frankfurt, Berlin and Hamburg).

These differences are also important for estimates of regional inequality among Germany districts. Cities play a large role: The average light intensity of rural areas is about 27 DN in the saturated data and 28 DN corrected data, but in urban areas its is about 66 DN vs. 79 DN. This translates into a spatial Gini coefficient of average lights per per km^2 of about 37.6 when the saturated data is used and about 42.3 once the corrected data is used. Note that increasing the size of the unit of observation means that we are discarding within district inequality in lights. Hence, this gap of approximately 5 Gini points is smaller than the average gap of the pixel level estimates but still sizable. Interestingly, once we consider per capita quantities, the discrepancy is much less pronounced. Now the spatial Gini coefficients of average per capita lights are 48.5 and 49.2 for the saturated and corrected data, respectively. This rather small difference may be owed to the fact that Germany is a decentralized territorial state, with only a comparatively small fraction of its population living in very big (bright) cities.

[To be completed later... We plan to use two different ways of isolating the economic significance of cities at the *global level* which is an important research question in the urban economics literature (e.g. see Storeygard, 2016). First, we will use the urban extents data produced by Schneider et al. (2010) and the MODIS (Moderate Resolution Imaging Spectroradiometer) satellites to classify urban regions. Most importantly, this data source is *independent* of the DMSP night lights. A second approach will just look at capital cities and cut out buffers (circles) around the city centroids corresponding to approximate city radius. As a byproduct of this exercise, we will produce two motivating graphs for this paper: in one we plot the intensity of top-coding in cities over the level of GDP per capita (suspecting that this will be a steep line), in the other we plot the growth rates of urban areas and non-urban areas over time using the original stable lights series, the true radiance-calibrated series and our Pareto imputed series.]

4.3 Regional and ethnic inequalities

[To be completed later... The main purpose of this application is to show that estimates of spatial and ethnic inequality can be very different when top coding is taken into account. We then show and discuss that this introduces bias *in favor* of the hypothesis of Alesina et al. (2016) when it comes to the relationship between inequality and underdevelopment (as richer countries will have more top-coding and hence appear to be more equal than they actually are). Using our corrected data and the radiance-calibrated data, we will also present a panel version of their cross-sectional base specification at the country-level.]

5 Concluding remarks

This paper deals with the problem of top-coding in satellite nighttime lights, which limits their use as a proxy for economic activity in studies of global development and regional convergence. When the full brightness of big cities is not measured, their continuing growth cannot be observed and their contribution to global economic activity is understated. This leads to an upward bias in estimates of regional convergence and a downward bias in estimates of spatial inequality.

We establish that top lights, just as top incomes, are Pareto distributed. On this basis, we then suggest a solution to the censoring problem based on methods from the top income literature: We augment the lower part of the saturated data with a Pareto tail. Our simple formula for the top-coding-corrected Gini coefficient in lights is computationally efficient because it does not require replacement of data at the pixel level. Our results for a global sample using 2% of all lit pixels show that top-coding correction makes a substantial difference. On average, it raises the spatial Gini coefficient by about 9 percentage points. These adjustments tend to be larger in countries which are (i) richer, (ii) smaller in size, or (iii) more urbanized, and *vice versa*.

For geo-referenced applications which require a completely new set of satellite images, we suggest a procedure for pixel level replacement based on simulated draws from a Pareto distribution. This method works well for Germany and we are currently implementing it on the global level. The goal is to produce a new top-coding corrected panel data set of night lights around the world from 1992 to 2013.

Our first applications to determine for which question in development economics and urban economics the top-coding problem is most severe open the door for further research. An intriguing finding is that the income elasticity of the saturated lights differs significantly between OECD and non-OECD countries but that of the radiance-calibrated lights does not. In line with this, we show that the light-output elasticity doubles in size and becomes significantly more robust after our correction. This challenges studies which argue that nighttime lights are not suited as proxy for income and output in developed economies. Overall, our results so far indicate that after appropriately accounting for top-coding, nighttime lights are a much better proxy of economic activity in all countries than previously assumed.

References

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

Anderson, G. (2004). Towards an Empirical Analysis of Polarization. *Journal of Econometrics 122*, 1–26.

Anderson, G., O. Linton, and T. W. Leo (2012). A Polarization-Cohesion Perspective on Cross-Country Convergence. *Journal of Economic Growth 17*, 49–69.

Atkinson, A., T. Piketty, and E. Saez (2011). Top incomes in the long run of history. *Journal of Economic Literature 49*, 3–71.

Atkinson, A. B. (2005). Top incomes in the UK over the 20th century. *Journal of the Royal Statistical Society: Series A (Statistics in Society) 168*(2), 325–343.

Bickenbach, F., E. Bode, P. Nunnenkamp, and M. Söder (2016). Night lights and regional gdp. *Review of World Economics 152*(2), 425–447.

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

Cirillo, P. (2013). Are your data really Pareto distributed. *Physica A: Statistical Mechanics and its Applications 392*, 5947–5962.

Cogneau, D. and Y. Dupraz (2014, June). Questionable Inference on the Power of Pre-Colonial Institutions in Africa. PSE Working Papers 2014-25, Paris School of Economics.

Cowell, F. (2013). Uk wealth inequality in international context. In J. Hills (Ed.), *Wealth in the UK: Distribution, Accumulation, and Policy*, pp. 35–62. Oxford University Press.

Dell, F. (2005). Top incomes in Germany and Switzerland over the twentieth century. *Journal of the European Economic Association 3*(2-3), 412–421.

Devarajan, S. (2013). Africa's statistical tragedy. *Review of Income and Wealth 59*, S9–S15.

Doll, C. N. (2008). Ciesin thematic guide to night-time light remote sensing and its applications. Technical report, Center for International Earth Science Information Network of Columbia University, Palisades, NY.

Duclos, J.-Y., J. Esteban, and D. Ray (2004). Polarization: Concepts, Measurement, Estimation. *Econometrica 72*, 1737–1772.

Elvidge, C., D. Ziskin, K. Baugh, B. Tuttle, T. Ghosh, D. Pack, E. Erwin, and M.Zhizhin (2009). A fifteen year record of global natural gas flaring derived from satellite data. *Energies 2*, 595–622.

Elvidge, C. D., F.-C. Hsu, K. E. Baugh, and T. Ghosh (2014). National trends in satellite-observed lighting. In Q. Weng (Ed.), *Global Urban Monitoring and Assessment through Earth Observation*, Taylor & Francis Series in Remote Sensing Applications, Chapter 6, pp. 97–120. CRC Press.

Esteban, J. and D. Ray (1994). On the Measurement of Polarization. *Econometrica 62*, 819–852.

Gabaix, X. (1999). Zipf's law for cities: An explanation. *Quarterly Journal of Economics 114*, 739–767.

Henderson, J., A. Storeygard, and D. Weil (2012). Measuring economic growth from outer space. *American Economic Review 102*, 994–1028.

Hodler, R. and P. Raschky (2014). Regional favoritism. *Quarterly Journal of Economics 129*, 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.

Lakner, C. and B. Milanovic (2015). Global income distribution: From the fall of the Berlin wall to the Great Recession. *World Bank Economic Review* (forthcoming).

Leßmann, C. and A. Seidel (2015). Regional inequality, convergence, and its determinants – a view from outer space. *CESifo Working Paper 5322*, 1–56.

Leßmann, C., A. Seidel, and A. Steinkraus (2015, December). Satellitendaten zur Schätzung von Regionaleinkommen – Das Beispiel Deutschland. *ifo Dresden berichtet 22*(06), 35–42.

Lorenz, M. O. (1905). Methods of measuring the concentration of wealth. *Publications of the American Statistical Association 9*(70), 209–219.

Mellander, C., J. Lobo, K. Stolarick, and Z. Matheson (2015). Night-time light data: A good proxy measure for economic activity? *PLoS ONE 10*, 1–18.

Michalopoulos, S. and E. Papaioannou (2013). Pre-colonial ethnic institutions and contemporary African development. *Econometrica 81*, 113–152.

Mookherjee, D. and A. Shorrocks (1982). A decomposition analysis of the trend in UK income inequality. *Economic Journal 92*(368), 886–902.

Newman, M. (2005). Power laws, pareto distributions and zipf's law. *Contemporary Physics 46*, 323–351.

Nordhaus, W. D. 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*, 217–246.

Pareto, V. (1897). La courbe de la répartition de la richesse. *Reprinted in Rivista di Politica Economica (1997) 87*, 647–700.

Piketty, T. (2003). Income inequality in France, 1901–1998. *Journal of Political Economy 111*(5), 1004–1042.

Pinkovskiy, M. and X. Sala-i Martin (2016). Lights, camera,... income! Estimating poverty using national accounts, survey means, and lights. *Quarterly Journal of Economics Forthcoming.*

Pisarenko, V. and M. Rodkin (2010). *Heavy-Tailed Distributions in Disaster Analysis*. Springer Science & Business Media.

Rozenfeld, H., D. Rybski, X. Gabaix, and H. Makse (2014). The area and population of cities: New insights from a different perspective on cities. *American Economic Review 101*, 2205–2225.

Schneider, A., M. A. Friedl, and D. Potere (2010). Mapping global urban areas using MODIS 500-m data: New methods and datasets based on 'urban ecoregions'. *Remote Sensing of Environment 114*(8), 1733 – 1746.

Small, C., C. Elvidge, D. Balk, and M. Montgomery (2012). Spatial scaling of stable night lights. *Remote Sensing of Environment 115*, 269–280.

Storeygard, A. (2016). Farther on down the road: transport costs, trade and urban growth in sub-Saharan Africa. *Review of Economic Studies* (forthcoming).

Weidmann, N. B. and S. Schutte (2016). Using night light emissions for the prediction of local wealth. *Journal of Peace Research*.

Appendix A A Pareto-augmented Gini coefficient for perfectly separated groups

Following Mookherjee and Shorrocks (1982) we begin by defining the Gini coefficient over multiple groups as

$$
\begin{aligned}
G &= \frac{1}{2N^2\mu} \sum_i \sum_j |y_i - y_j| && \text{(A-1)}\\
&= \frac{1}{2N^2\mu} \sum_k \left(\sum_{i \in N_k} \sum_{j \in N_k} |y_i - y_j| + \sum_{i \in N_k} \sum_{j \notin N_k} |y_i - y_j| \right) && \text{(A-2)}\\
&= \sum_k \left(\frac{N_k}{N} \right)^2 \frac{\mu_k}{\mu} G_k + \frac{1}{2N^2\mu} \sum_k \sum_{i \in N_k} \sum_{j \notin N_k} |y_i - y_j| \, . && \text{(A-3)}
\end{aligned}
$$

where G_K is the within group Gini coefficient of group k. The second term is a measure of group overlap including their between group differences.

Perfect separation (no overlap between groups) implies $\sum_{i \in N_k} \sum_{j \in N_h} |y_i - y_j| = N_k N_h |\mu_k - \mu_h|$. Hence, we can simplify equation (A-3) to

$$
G = \sum_k \left(\frac{N_k}{N} \right)^2 \frac{\mu_k}{\mu} G_k + \sum_k \sum_h \frac{N_k N_h}{2N^2\mu} |\mu_k - \mu_h| \, . \tag{A-4}
$$

With two bottom and top groups $k, h \in \{B, T\}$ (where $\mu_T > \mu_B$) and some algebra, this becomes

$$
G = \left(\frac{N_B}{N} \right)^2 \frac{\mu_B}{\mu} G_B + \left(\frac{N_T}{N} \right)^2 \frac{\mu_T}{\mu} G_T + \left[\left(\frac{N_T}{N} \right)^2 \frac{\mu_T}{\mu} - \frac{N_T}{N} \right] . \tag{A-5}
$$

Now define the population (pixel) shares ω_B and ω_T, where $\omega_T = 1 - \omega_B$ and the group's share of all income (light) as $\phi_B = \omega_B \frac{\mu_B}{\mu}$ and $\phi_T = \omega_T \frac{\mu_T}{\mu}$. The formula

$$
G = \omega_B \cdot \phi_B \cdot G_B + \omega_T \cdot \phi_T \cdot G_B + [\phi_T - \omega_T] \tag{A-6}
$$

as presented by Cowell (2013) for the top-income case, can be simplified by using the Pareto distribution for the top share, whose Gini is $G_T = \frac{1}{2\alpha - 1}$:

$$
G = \omega_B \cdot \phi_B \cdot G_B + (1 - \omega_B) \cdot \phi_T \cdot \frac{1}{2\alpha - 1} + [\phi_T - (1 - \omega_B)] \tag{A-7}
$$

Plugging in $\phi_B = \omega_B \frac{\mu_B}{\mu}$, $\phi_T = \omega_T \frac{\mu_T}{\mu}$, $\mu = \omega_B \mu_B + \omega_T \mu_T$ as well as the mean of the Pareto-distributed top $\mu_T = \frac{\alpha}{\alpha - 1} y_c$ yields eq. (8) from the text.

Appendix B   Additional Tables and Figures

Table A-1: Top-coding correction per country in 2010, Part 1

Country	Gini Unadj.	Top share ω_T	Mean μ_B	Gini G_B	Gini adj.	Diff.
AFG	0.4286	0.1116	18.80	0.3967	0.5603	0.1318
AGO	0.4130	0.2185	19.66	0.3924	0.5926	0.1796
ALA	0.3240	0.0477	18.73	0.2859	0.4155	0.0915
ALB	0.3665	0.0142	12.19	0.3419	0.4108	0.0443
ARE	0.4279	0.1939	19.68	0.4170	0.5950	0.1671
ARG	0.4383	0.0700	15.22	0.3877	0.5497	0.1114
ARM	0.3671	0.0448	16.11	0.3211	0.4553	0.0882
ATA	0.1073	0.0000	6.98	0.1073	0.1073	0.0000
ATF	0.0516	0.0000	4.34	0.0516	0.0516	0.0000
AUS	0.4900	0.0742	11.49	0.4084	0.6117	0.1217
AUT	0.3343	0.0281	17.46	0.3077	0.3957	0.0614
AZE	0.3988	0.0296	11.61	0.3487	0.4787	0.0799
BDI	0.4063	0.0552	17.63	0.3780	0.5004	0.0940
BEL	0.2409	0.2450	32.69	0.2374	0.4490	0.2081
BEN	0.4062	0.0377	13.80	0.3634	0.4897	0.0835
BFA	0.4188	0.1068	15.92	0.3440	0.5614	0.1426
BGD	0.3531	0.0143	8.87	0.3090	0.4117	0.0586
BGR	0.3387	0.0098	11.10	0.3166	0.3745	0.0359
BHR	0.1572	0.6341	33.67	0.2682	0.4676	0.3104
BHS	0.5234	0.0970	12.03	0.4586	0.6507	0.1273
BIH	0.3214	0.0154	16.15	0.3035	0.3615	0.0401
BLR	0.3768	0.0396	15.79	0.3380	0.4586	0.0818
BLZ	0.3498	0.0071	11.29	0.3356	0.3765	0.0267
BOL	0.4516	0.0826	14.06	0.3834	0.5765	0.1249
BRA	0.4139	0.0615	15.13	0.3592	0.5212	0.1073
BRN	0.3411	0.2198	25.42	0.3353	0.5270	0.1859
BTN	0.3416	0.0032	13.35	0.3372	0.3528	0.0113
BWA	0.4286	0.0265	13.00	0.3966	0.4919	0.0634
CAF	0.3582	0.0000	13.19	0.3582	0.3582	0.0000
CAN	0.4020	0.0353	13.19	0.3551	0.4835	0.0815
CHE	0.3159	0.0798	24.93	0.2951	0.4264	0.1105
CHL	0.4296	0.0741	15.42	0.3749	0.5460	0.1164
CHN	0.4258	0.0591	15.11	0.3791	0.5296	0.1037
CIV	0.3566	0.0182	8.85	0.3010	0.4280	0.0714
CMR	0.4151	0.0678	13.62	0.3382	0.5369	0.1217
COD	0.4018	0.0909	17.03	0.3411	0.5313	0.1294
COG	0.4351	0.1165	17.13	0.3871	0.5757	0.1405
COL	0.4121	0.0349	12.10	0.3594	0.4970	0.0848
COM	0.2583	0.0000	7.10	0.2583	0.2583	0.0000
CPV	0.4040	0.0132	10.06	0.3770	0.4511	0.0471
CRI	0.3772	0.0250	10.32	0.3199	0.4558	0.0786
CUB	0.4062	0.0217	10.26	0.3619	0.4746	0.0684
CYP	0.4180	0.0688	16.11	0.3690	0.5282	0.1103
CZE	0.3219	0.0324	19.93	0.2994	0.3855	0.0636
DEU	0.3190	0.0612	21.90	0.2880	0.4176	0.0986
DJI	0.4655	0.1992	15.21	0.4093	0.6379	0.1724
DMA	0.2309	0.0000	7.56	0.2309	0.2309	0.0000
DNK	0.3346	0.0525	19.09	0.2960	0.4290	0.0943
DOM	0.4463	0.0348	10.37	0.3868	0.5360	0.0897
DZA	0.4038	0.0355	14.36	0.3661	0.4813	0.0774
ECU	0.4062	0.0435	13.15	0.3501	0.5012	0.0950
EGY	0.3578	0.2634	23.99	0.3740	0.5571	0.1993
ERI	0.3894	0.0207	11.14	0.3517	0.4540	0.0646
ESH	0.4440	0.0762	16.77	0.4129	0.5564	0.1124
ESP	0.4285	0.0709	15.95	0.3820	0.5402	0.1117
EST	0.3668	0.0792	19.82	0.3275	0.4836	0.1168
ETH	0.3815	0.0307	12.61	0.3342	0.4611	0.0797

Table A-2: Top-coding correction per country in 2010, Part 2

Country	Gini Unadj.	Top share ω_T	Mean μ_B	Gini G_B	Gini adj.	Diff.
FIN	0.3720	0.0713	20.70	0.3450	0.4774	0.1054
FJI	0.3033	0.0000	9.48	0.3033	0.3033	0.0000
FLK	0.2699	0.0000	7.95	0.2699	0.2699	0.0000
FRA	0.3652	0.0531	17.70	0.3250	0.4608	0.0956
FRO	0.2685	0.0000	22.16	0.2685	0.2685	0.0000
FSM	0.1981	0.0000	6.39	0.1981	0.1981	0.0000
GAB	0.4309	0.0617	13.68	0.3710	0.5446	0.1137
GBR	0.3817	0.1160	20.35	0.3411	0.5228	0.1411
GEO	0.4020	0.0638	17.12	0.3626	0.5076	0.1056
GHA	0.4398	0.0522	10.74	0.3511	0.5536	0.1138
GIN	0.3592	0.0159	16.35	0.3456	0.3991	0.0399
GLP	0.3952	0.0870	18.56	0.3516	0.5174	0.1222
GMB	0.4083	0.0706	19.30	0.3852	0.5123	0.1040
GNB	0.3051	0.0000	19.06	0.3051	0.3051	0.0000
GNQ	0.2944	0.1000	31.01	0.3053	0.4158	0.1214
GRC	0.4111	0.0310	13.08	0.3720	0.4852	0.0740
GRL	0.3533	0.0000	14.71	0.3533	0.3533	0.0000
GTM	0.3690	0.0204	10.43	0.3222	0.4367	0.0677
GUF	0.4213	0.0617	14.78	0.3696	0.5315	0.1102
GUM	0.3116	0.2250	28.39	0.3155	0.4994	0.1878
GUY	0.3510	0.0198	12.80	0.3183	0.4093	0.0583
HKG	0.2005	0.5000	33.06	0.2654	0.4814	0.2809
HND	0.3885	0.0187	10.87	0.3518	0.4483	0.0598
HRV	0.3376	0.0380	18.20	0.3059	0.4133	0.0757
HTI	0.4235	0.0345	13.95	0.3917	0.5010	0.0775
HUN	0.3444	0.0324	15.47	0.3055	0.4192	0.0748
IDN	0.3859	0.0318	11.70	0.3291	0.4705	0.0846
IMN	0.3417	0.0000	16.43	0.3417	0.3417	0.0000
IND	0.3549	0.0174	10.85	0.3158	0.4135	0.0585
IRL	0.3472	0.0238	13.79	0.3111	0.4104	0.0632
IRN	0.4324	0.0432	13.17	0.3833	0.5224	0.0899
IRQ	0.4142	0.0540	14.95	0.3645	0.5133	0.0992
ISL	0.3896	0.0889	20.45	0.3592	0.5074	0.1178
ISR	0.3916	0.2471	21.40	0.3928	0.5811	0.1896
ITA	0.3668	0.0985	21.83	0.3390	0.4941	0.1272
JAM	0.3732	0.0213	11.87	0.3338	0.4354	0.0622
JOR	0.4136	0.1124	19.07	0.3777	0.5493	0.1357
JPN	0.4470	0.1129	16.57	0.3998	0.5845	0.1375
KAZ	0.3983	0.0396	14.99	0.3592	0.4805	0.0823
KEN	0.3788	0.0331	11.91	0.3218	0.4674	0.0886
KGZ	0.3466	0.0166	12.27	0.3153	0.3981	0.0515
KHM	0.4118	0.0927	17.93	0.3670	0.5377	0.1260
KIR	0.1413	0.0000	5.75	0.1413	0.1413	0.0000
KOR	0.4066	0.1306	19.29	0.3645	0.5538	0.1472
KWT	0.4669	0.1471	16.01	0.4172	0.6167	0.1497
LAO	0.3680	0.0779	20.62	0.3347	0.4798	0.1118
LBN	0.3568	0.0426	22.40	0.3423	0.4251	0.0683
LBR	0.2740	0.0000	17.69	0.2740	0.2740	0.0000
LBY	0.4197	0.1186	17.87	0.3695	0.5618	0.1421
LCA	0.3791	0.0435	18.32	0.3550	0.4598	0.0806
LKA	0.2992	0.0099	10.22	0.2713	0.3409	0.0416
LSO	0.4251	0.0534	13.50	0.3686	0.5304	0.1052
LTU	0.3630	0.0455	18.31	0.3314	0.4467	0.0836
LUX	0.2490	0.1243	28.62	0.2170	0.4001	0.1511
LVA	0.3680	0.0746	21.25	0.3409	0.4748	0.1068
MAR	0.4357	0.0539	13.43	0.3793	0.5385	0.1028
MDA	0.3344	0.0114	11.07	0.3091	0.3768	0.0424
MDG	0.4204	0.0302	12.75	0.3813	0.4930	0.0726
MEX	0.4523	0.0429	11.16	0.3879	0.5481	0.0958
MKD	0.3727	0.0159	11.73	0.3434	0.4223	0.0496
MLI	0.3956	0.0945	17.95	0.3409	0.5257	0.1301

Table A-3: Top-coding correction per country in 2010, Part 3

Country	Gini Unadj.	Top share ω_T	Mean μ_B	Gini G_B	Gini adj.	Diff.
MMR	0.3813	0.0360	13.58	0.3324	0.4653	0.0840
MNE	0.3496	0.0250	18.25	0.3311	0.4028	0.0532
MNG	0.4126	0.0434	15.63	0.3768	0.4954	0.0827
MOZ	0.4375	0.0962	15.76	0.3796	0.5698	0.1323
MRT	0.4737	0.0597	13.25	0.4222	0.5760	0.1023
MTQ	0.3171	0.2069	28.50	0.3246	0.4963	0.1793
MUS	0.3963	0.0600	16.00	0.3451	0.5012	0.1049
MWI	0.3971	0.0550	15.26	0.3458	0.5003	0.1032
MYS	0.4234	0.1019	16.40	0.3608	0.5592	0.1358
NAM	0.4488	0.0259	10.39	0.4036	0.5198	0.0710
NCL	0.4294	0.0368	9.96	0.3545	0.5268	0.0974
NER	0.4249	0.0274	11.65	0.3820	0.4957	0.0708
NGA	0.4146	0.0460	13.54	0.3612	0.5099	0.0954
NIC	0.3991	0.0279	11.49	0.3502	0.4754	0.0763
NLD	0.2696	0.1780	29.57	0.2477	0.4459	0.1763
NOR	0.3483	0.0961	22.98	0.3215	0.4731	0.1248
NPL	0.2987	0.0095	9.58	0.2693	0.3409	0.0423
NZL	0.4584	0.0410	11.10	0.3990	0.5512	0.0928
OMN	0.4205	0.1182	19.50	0.3924	0.5565	0.1360
PAK	0.3490	0.0161	11.30	0.3144	0.4023	0.0533
PAN	0.4268	0.0412	13.44	0.3814	0.5139	0.0870
PER	0.4314	0.0519	12.99	0.3704	0.5343	0.1029
PHL	0.4110	0.0389	11.24	0.3424	0.5071	0.0961
PNG	0.3758	0.0167	11.64	0.3454	0.4281	0.0523
POL	0.2996	0.0430	21.40	0.2718	0.3779	0.0783
PRI	0.2994	0.2084	28.33	0.2916	0.4861	0.1867
PRK	0.3636	0.0083	10.45	0.3453	0.3960	0.0325
PRT	0.4122	0.0846	17.98	0.3729	0.5320	0.1198
PRY	0.4155	0.0766	14.92	0.3458	0.5389	0.1234
PSE	0.3500	0.1516	24.90	0.3367	0.5056	0.1556
PYF	0.4111	0.0213	9.46	0.3653	0.4846	0.0736
QAT	0.4528	0.1596	16.68	0.4032	0.6097	0.1569
REU	0.4653	0.1319	17.39	0.4457	0.6069	0.1417
ROU	0.3331	0.0186	12.69	0.2993	0.3901	0.0570
RUS	0.3990	0.0339	13.75	0.3585	0.4771	0.0781
RWA	0.4162	0.0771	16.48	0.3695	0.5358	0.1196
SAU	0.4457	0.1038	16.21	0.3934	0.5783	0.1326
SDN	0.4359	0.0645	13.64	0.3696	0.5487	0.1128
SEN	0.4398	0.0492	12.60	0.3805	0.5396	0.0997
SGP	0.0277	0.8974	47.75	0.0877	0.4065	0.3787
SGS	0.0000	0.0000	5.00	0.0000	0.0000	0.0000
SLB	0.3024	0.0000	8.45	0.3024	0.3024	0.0000
SLE	0.3821	0.0280	19.50	0.3718	0.4369	0.0547
SLV	0.3718	0.0185	10.72	0.3310	0.4321	0.0603
SOM	0.3615	0.0060	12.41	0.3520	0.3820	0.0205
SRB	0.3367	0.0256	16.74	0.3108	0.3958	0.0592
STP	0.3006	0.0000	9.77	0.3006	0.3006	0.0000
SUR	0.3950	0.0791	16.81	0.3373	0.5170	0.1220
SVK	0.3222	0.0185	16.25	0.3015	0.3695	0.0473
SVN	0.3278	0.0170	16.82	0.3102	0.3697	0.0419
SWE	0.3712	0.0542	18.47	0.3368	0.4646	0.0934
SWZ	0.3430	0.0000	10.80	0.3430	0.3430	0.0000
SYR	0.3897	0.0264	13.21	0.3530	0.4563	0.0666
TCA	0.4366	0.0000	21.23	0.4366	0.4366	0.0000
TCD	0.4402	0.0718	16.50	0.4028	0.5480	0.1078
TGO	0.4302	0.0556	11.64	0.3489	0.5464	0.1162
THA	0.3890	0.0680	16.49	0.3330	0.5020	0.1130
TJK	0.3152	0.0098	9.21	0.2850	0.3596	0.0444
TKM	0.4181	0.0394	14.39	0.3794	0.4996	0.0815
TLS	0.3770	0.0084	10.33	0.3599	0.4099	0.0329
TON	0.2703	0.0000	7.81	0.2703	0.2703	0.0000

Table A-4: Top-coding correction per country in 2010, Part 4

Country	Gini Unadj.	Top share ω_T	Mean μ_B	Gini G_B	Gini adj.	Diff.
TTO	0.3355	0.2678	24.04	0.3136	0.5421	0.2066
TUN	0.4366	0.0371	11.55	0.3814	0.5243	0.0876
TUR	0.4100	0.0317	12.04	0.3619	0.4901	0.0801
TWN	0.3825	0.2563	22.20	0.3978	0.5766	0.1941
TZA	0.3985	0.0319	12.65	0.3508	0.4765	0.0780
UGA	0.3940	0.0538	13.85	0.3282	0.5020	0.1079
UKR	0.3555	0.0172	11.33	0.3205	0.4125	0.0570
URY	0.4602	0.0589	14.08	0.4149	0.5616	0.1014
USA	0.4490	0.0849	14.46	0.3834	0.5744	0.1255
UZB	0.3788	0.0218	12.50	0.3435	0.4395	0.0607
VEN	0.4319	0.0503	12.63	0.3691	0.5346	0.1027
VNM	0.3900	0.0332	14.16	0.3504	0.4658	0.0758
VUT	0.4183	0.0195	16.85	0.4086	0.4609	0.0427
WSM	0.3165	0.0000	7.02	0.3165	0.3165	0.0000
YEM	0.4131	0.0293	11.73	0.3667	0.4895	0.0764
ZAF	0.4605	0.0590	12.24	0.3921	0.5697	0.1092
ZMB	0.4049	0.0732	16.15	0.3482	0.5219	0.1171
ZWE	0.4345	0.0302	11.07	0.3870	0.5140	0.0795

Notes: The table reports the results from the top-coding correction in all countries. We report the unadjusted Gini coefficient first, then provide statistics of the top share of pixels, the bottom mean and Gini coefficient, and the new composite Gini coefficient after applying eq. (8) as presented in the text. The last column demonstrates how sizable the correction can be by computing the difference in the two Gini coefficients. All estimates use $\alpha = 1.75$ and $y_c = 55$ for the Pareto-imputation.

Appendix C Between and within satellite measurement errors

Apart from top-coding, there is another inherent limitation of the DMSP-OLS satellite system that makes it difficult to compare its images across time (but does not affect comparisons across space within one image). This limitation arises because the pictures are recorded using a variable gain setting. The sensor gain basically works like a pre-amplifier: it needs to be high if the satellites are supposed to register very dim lights and low if they are to pick up very bright lights. Since the gain of the DMSP system is variable and the value is not recorded on board, it cannot be recovered or linked back to a physical quantity like radiance (Elvidge et al., 2009; Doll, 2008). This problem is only made worse by the fact that different satellites have sensors that deliver, by their different construction, brighter or dimmer pictures and those sensors tend to degrade over time. As a result, the stable lights series suffers from jumps in the time-series dimension that are caused both by switches in the satellite delivering the images (between satellite measurement error) and changes in the ability of the sensors to detect light over time (within satellite measurement error). Figure A-1 plots the average light intensity in Sicily as obtained from each separate satellite and shows how severe these jumps can be.

Figure A-1: Average lights in Sicily according to each satellite

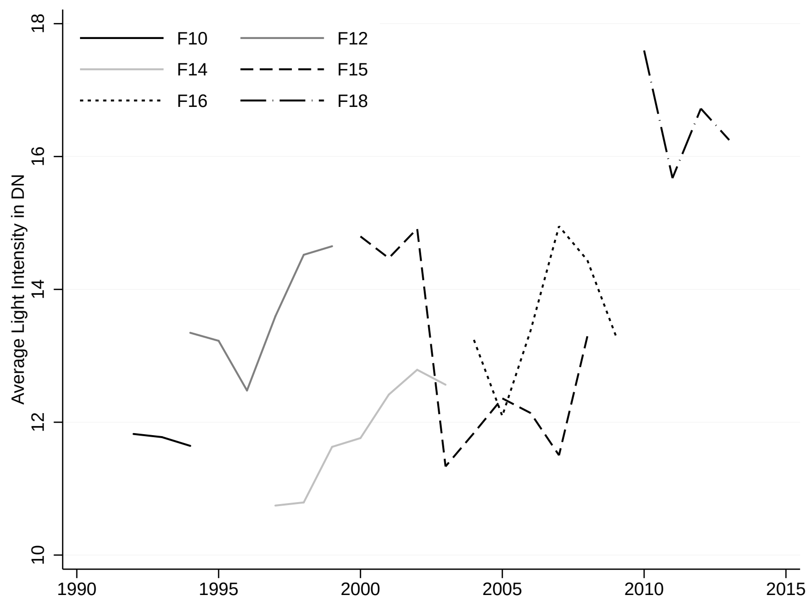

To be clear, the nature of the between and within satellite errors is such that there are only three possible types of perturbations to the data: a) the brightness of the entire intermediate range ($y \in [1, 62]$) is shifted by a constant in each image, b) the extent of top-coding increases when the gain is higher on average than before, and c) the gain setting is lower on average than before, leading to more bottom-coding; that is, a decrease in the ability of the satellites to pick up dim lights. NOAA then applies a series of filters to the data to remove background noise, ephemeral lights and more, but these adjustments are uniformly applied to all images that make up a series of so-called composites.

Economists usually deal with these measurement errors by including time fixed effects in the regression of interest (e.g. see Henderson et al., 2012; Michalopoulos and Papaioannou, 2013; Hodler and Raschky, 2014). Chen and Nordhaus (2011) and Henderson et al. (2012) provide a detailed discussion and estimates of measurement errors in lights and in GDP but do not correct the underlying data. The problem with using time fixed-effects is that these are only a valid remedy in panel studies where the light data is used in conjunction with other variables. If the time-series properties of the light data itself are of interest, then some form of adjustment needs to be undertaken to smooth out the artificial jumps in the series.

The producers of the lights data at NOAA suggest a very different procedure. Elvidge et al. (2009) propose to "inter-calibrate" each image by scaling it to match the brightest image within a fixed reference area. Specifically, they argue that Sicily covers the entire dynamic spectrum of the saturated data and experienced little change in lighting since 1992. Then, they run quadratic regressions of the form $E[\text{F12 in 1999}|X] = \beta_0 + \beta_1 X + \beta_2 X^2$ where X stands for the corresponding pixel from any of the other satellites. The estimated coefficients can then be applied to re-scale the global images and recalculate all statistics of interest. Chen and Nordhaus (2011) already noticed that this procedure is a bit awkward in the sense that does not impose any useful parameter restrictions; it allows negative estimates of the intercept and, more generally, often produces estimates outside of the observed data range. Nevertheless, this method has not been systematically analyzed so far (apart from the original results presented in Elvidge et al., 2009, 2014).

To reproduce their approach, we isolate Sicily using the GADM (Global Administrative Boundaries Dataset) and then build a pixel-level data set of all 35 satellites.[20] We then run regressions of the reference satellite (F12 in 1999) on each image and predict the adjusted value for each pixel. Table A-5 shows the results of this exercise. As expected, the R^2 is generally high and exceeds 0.90 for all but four satellite-years. For F16 in 2009 to F18 in 2011, it falls substantially below 0.90. It is not exactly clear if this is in part due to slight displacement of the pixels,[21] if this is purely due to sensor differences, or if this could also be the product of *genuine economic effects* (such as the impact of the Great Recession in Europe since 2008).

Figure A-2 shows why this approach is very problematic. It plots the average light

[20] Note that we also align the underlying pixel grid, so that each pixel is matched to its nearest neighbor across various images. Elvidge et al. (2009) first project the data into a Mollweide equal area projection and then proceed with the analysis. Since Sicily is very small, these differences are likely to be immaterial. However, it does ensure that all of our pixels are matched and N is the same across the panel which highlights the properties of the method much more clearly.

[21] The exact location can vary between one and two kilometers. These distortions are introduced during the compositing process undertaken at both on board of the satellites and at NOAA.

Figure A-2: Results of "inter-satellite calibration" and real GDP in Sicily

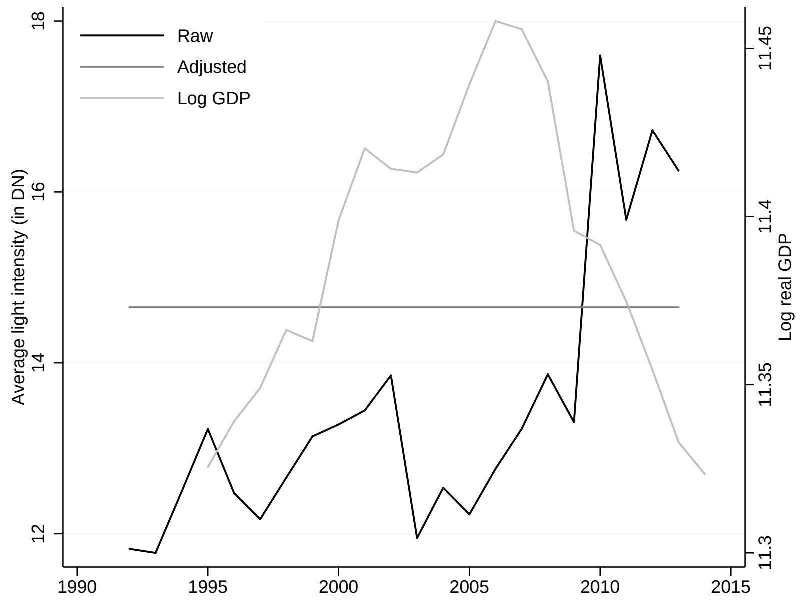

intensity in Sicily before and after the adjustment. We simply average the data whenever we have more than one satellite at our disposal. One the one hand, the adjusted average light intensity is clearly perfectly stable across all years. In fact, this follows from a basic property of OLS regression; namely, the line always passes through the mass point $\{\bar{Y}, \bar{X}\}$. Note that this property also implies that the sum of light will be stable as well, if it is estimated on the same sample (since $N \cdot \text{E}[Y] = N \cdot \text{E}[Y|\bar{X}]$). On the other hand, there is little reason to assume that average (or total) lights in Sicily are actually that stable. In fact, there is lots of evidence to the contrary. As Figure A-2 also shows, regional national accounts indicate that real GDP in Sicily grew substantially over the period form 1995 to 2007 and then fell again below its initial value by 2014. Likewise, electricity consumption has increased steadily over the entire period. The raw data also shows that the F18 satellite is actually brighter on average than its antecedents (but its data was not available when Elvidge et al., 2009, proposed to scale to the F12 satellite).

Are there sensible alternatives? There are two ways to remedy this situation and still produce a reliable time-series of night time lights. A first option is to find a reference point where lights can actually reasonably be assumed to be constant over the two decades in question. We are not too hopeful that such a place exists. A second and better option is use a different calibration approach. One promising avenue is to exploit the fact that we have overlapping satellite-years for all but the last switch of satellites (F16 to F18). We could therefore run panel regressions of all satellites on a set of satellite dummies (to account for linear shifts), satellite time trends (to account for sensor degradation) and year dummies (or a linear trend). This approach should in theory be able to separately identify the between satellite differences, the within satellite time trends, while absorbing the remaining time-series variation in the time dummies (or trend) for the series from 1992 to 2009. Later versions of this paper will include a calibrated series.

Table A-5: "Inter-satellite calibration" regression coefficients

Satellite	Year	β_0	β_1	β_2	R^2	N
F10	1992	0.2364	1.3649	-0.0055	0.900	37887
F10	1993	-1.6439	1.6338	-0.0097	0.936	37887
F10	1994	0.4946	1.3927	-0.0064	0.929	37887
F12	1994	1.1103	1.0156	-0.0000	0.916	37887
F12	1995	-0.0835	1.2111	-0.0034	0.925	37887
F12	1996	0.7600	1.1903	-0.0027	0.936	37887
F12	1997	-0.2448	1.1572	-0.0022	0.931	37887
F12	1998	-0.2424	1.0588	-0.0011	0.956	37887
F12	1999	0.0000	1.0000	0.0000	1.000	37887
F14	1997	-0.6512	1.6913	-0.0108	0.916	37887
F14	1998	0.2655	1.5840	-0.0093	0.969	37887
F14	1999	-0.8969	1.5694	-0.0087	0.970	37887
F14	2000	0.6693	1.3498	-0.0057	0.935	37887
F14	2001	-0.1938	1.3484	-0.0055	0.945	37887
F14	2002	0.8926	1.1701	-0.0032	0.929	37887
F14	2003	-0.1146	1.3156	-0.0050	0.944	37887
F15	2000	-1.1409	1.1311	-0.0022	0.940	37887
F15	2001	-1.0157	1.1246	-0.0015	0.959	37887
F15	2002	-0.0350	0.9547	0.0010	0.964	37887
F15	2003	-0.4731	1.5599	-0.0087	0.934	37887
F15	2004	0.7588	1.3035	-0.0047	0.949	37887
F15	2005	-0.2145	1.3421	-0.0051	0.939	37887
F15	2006	0.1245	1.3311	-0.0049	0.942	37887
F15	2007	1.2463	1.2801	-0.0042	0.910	37887
F15	2008	3.5115	0.7306	0.0032	0.916	37887
F16	2004	0.3563	1.1620	-0.0029	0.919	37887
F16	2005	-0.8824	1.4756	-0.0072	0.940	37887
F16	2006	0.1760	1.1191	-0.0013	0.926	37887
F16	2007	0.3880	0.9136	0.0013	0.949	37887
F16	2008	0.2815	0.9973	-0.0001	0.946	37887
F16	2009	2.3508	0.9401	-0.0005	0.807	37887
F18	2010	1.8984	0.5306	0.0060	0.839	37887
F18	2011	2.3274	0.7302	0.0017	0.755	37887
F18	2012	1.0646	0.6666	0.0045	0.939	37887
F18	2013	1.0978	0.7354	0.0030	0.939	37887

Notes: The table shows results using a quadratic regression of the form: $E[\text{F12 in 1999}|X] = \beta_0 + \beta_1 X + \beta_2 X^2$ where X stands for the corresponding pixel from any of the other satellites. The data are are the stable lights data at the pixel level after applying the GADM boundaries to isolate the island of Sicily. The data are adjusted such that the cells of the satellite images align, since F162009, F182010 and F182011 are shifted by half a pixel for unknown reasons. Each grid cell is 30 arc seconds by 30 arc seconds. We do not project the data onto an equal area grid to not induce distortions in the projection and interpolation process. Instead, we weight each pixel by its land area for the applications.