Krause, Melanie; Bluhm, Richard

Conference Paper

Top Lights - Bright Spots and their Contribution to Economic Development

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2016: Demographischer Wandel - Session: Development Economics: Nighttime Luminosity, No. E13-V1

Provided in Cooperation with:
Verein für Socialpolitik / German Economic Association

Top Lights

Bright Spots and their Contribution to Economic Development

Very Preliminary. Please do not cite or distribute.
February 29, 2016

Abstract

Satellite data on nighttime luminosity is an increasingly popular proxy for economic activity in developing countries. However, their use for analyzing inequality and convergence on a global scale is severely limited by top-coding of the NOAA satellite images, which fail to accurately capture the brightness of large and densely populated cities, as well as by comparability problems between satellites. As a result, they severely underestimate differences between urban and rural regions, and developed and developing countries. We propose a new and easy-to-use procedure to correct for top-coding in nighttime lights, which borrows from the top incomes literature. We show that just as top incomes, top lights are Pareto distributed. We then derive simple formulas for the top-coding adjusted spatial Gini coefficient and top-coding adjusted average light intensity. Using data for Germany we show that by correcting for top-coding of the top 2% of lights, we can account for up to 40% of the difference between saturated and unsaturated satellites. We also analyze corrections for between and within satellite measurement errors. Finally, we present three economic applications to determine where the influence of top-coding is most severe. We show that top-coding and satellite calibration affects estimates of the income elasticity of light, regional inequalities and urban-rural differences.

JEL Classification: D3, O1, O18, C34
Keywords: Development, Inequality, Nighttime Lights, Top-Coding, Top Incomes

1 Introduction

Economic activity in countries with low quality statistical data is hard to track. This is a problem which not only plagues studies focusing on developing countries but also those dealing with growth, inequality and development on a global scale. Advanced economies tend to be over-represented in empirical studies, while less developed countries, particularly in Africa, are left out due to data unavailability. This data constraint is so severe that it has been referred to as a 'statistical tragedy' (Devarajan, 2013).

Data from weather satellites circling the earth at night while capturing light emissions are increasingly seen as a way out of this dilemma. A growing literature in economics is now using 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; Dreher and Lohmann, 2015). The advantages of the night lights data are obvious. The data are publicly available as a time series from 1992 onwards for nearly all parts of the world below the Arctic circle. They have a high resolution compared to regional national accounts data and they are measured uniformly across the globe. Hence, they are deemed to be comparable both within and between countries. As Henderson et al.'s (2012) seminal paper shows, night lights predict output growth 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. Top-coding rises with income and hence distorts estimates of regional inequality and convergence. The main contribution of this paper is to demonstrate that the upper tail of the distribution of night light intensities follows a Pareto law and to present top-coding corrected estimates of average lights and spatial inequalities.

Geo-referenced images of night lights are typically obtained from the National Geophysical Data Center (NGDC) at the National Oceanic Administration Agency (NOAA), whose DMSP-OLS satellites 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 preprocessed (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 a data set of annual light intensities from 1992 to 2013 at a resolution of 30 by 30 arc seconds for every pixel around the globe[1] (corresponding approximately to 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 at the turn of the millennium.

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

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

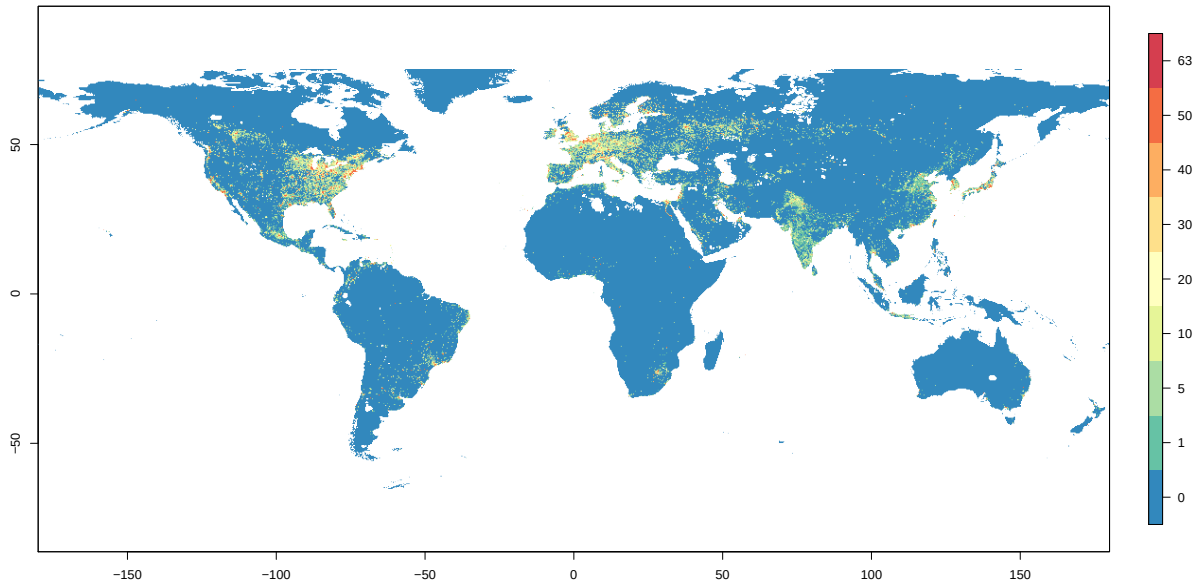

The values of this so-called 'stable lights'[2] product are recorded in a fixed range of digital numbers (DN) from 0 (missing or completely dark) to 63 (bright). Due to sensor saturation, the satellites are not able to capture a light intensity higher than 63, rendering them unable to distinguish between a mid-sized city and a booming metropolis in most developed countries. This top-coding of the data understates the differences in lights between rural and bright urban areas, inducing downward bias in inequality measures and upward bias in the speed of convergence. It is an underestimated 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), and estimates of spatial inequalities (Alesina et al., 2016).

Globally, not too many pixels are top-coded, but the truncation of the scale is enormous. Cities like Berlin, New York, or Tokyo are more than ten times brighter than recorded by the stable lights data. As an example, consider Figure 2 for Germany in 1999, where the saturated data on the right hit the 63 DN threshold in most big urban agglomeration. The saturated data on the left do not differentiate among larger cities. The non-saturated lights on the right clearly allow us to locate the brightest spots (Berlin, Hamburg, Munich) and contrast them to dimmer cities. Top-coding tends to affect developed countries more than their less-developed counterparts, but 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 in part explain why night lights are more strongly correlated with economic variables in developing rather than developed countries (see Pinkovskiy and Sala-i Martin, 2014; Nordhaus and Chen, 2015).

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

This conjecture was recently confirmed by Mellander et al. (2015) for Sweden: Using high resolution grids of 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 later show that this point carries over to estimates of the income elasticity of lights in OECD countries.

Figure 2: Germany, Satellite F12 in 1999, saturated (stable-lights) and unsaturated

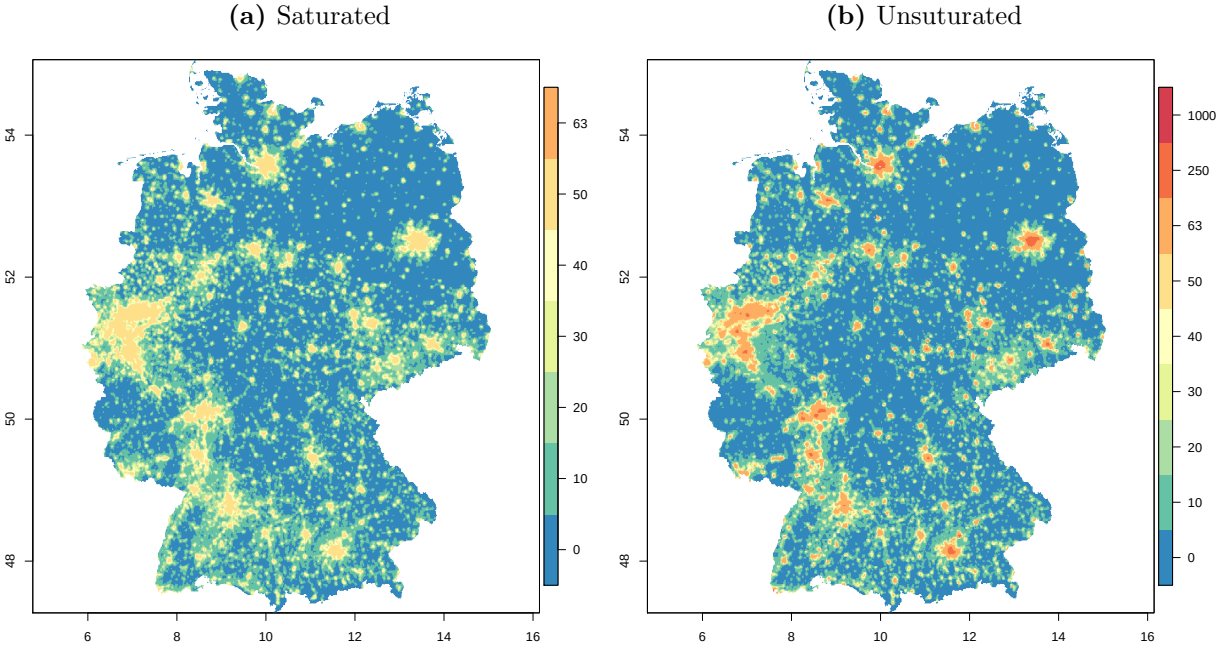

Top-coding has received little attention in the literature so far since we lacked reliable time-series data on non-saturated lights.[3] This situation has now changed. For seven years, additional satellites flew with various fixed gain settings that are less sensitive to light and thus capable of capturing the upper part of the light distribution. The resulting 'radiance-calibrated' data are not top-coded but at least three major problems remain:[4] (i) The data are only available for a few years whereas 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 data series. (iii) The radiance-calibrated series exhibits great variability over the years and is not strictly comparable across images from different satellites and years.

Measurement errors are also present in the 'stable lights' data and severely limit their value as a panel data set. 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

[3] Until recently only one cross-section of unsaturated lights in 2006 was made available by the NGDC. Henderson et al. (2012) make use of this data to gauge the influence of top-coding, but are not able to replace the stable lights data with this product which is a much more stringent robustness check.

[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.

across and *within* satellites. In settings such as panel regressions, we may resort to a combination of satellite fixed effects and time fixed effects, but in others, it is not possible to ensure the comparability of the DNs in the time-series dimension. 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 (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 and non-reporting, 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 combine the two data sources and suggest a three step approach of correcting the lights data. First, we estimate the Pareto parameter α using the radiance-calibrated data, then we transfer these estimates to the corresponding years of the saturated data, interpolating the intermediate years, and finally we combine the lower part of the observed distribution with the upper part of the Pareto distribution. This approach can even be extended to correct the data at the *pixel level* and we are currently working on implementing such a correction.

First results show that our top-coding correction makes a substantial difference. The estimated Gini coefficient of lights in Germany increases by 2 to 5 points depending on the year. This difference accounts for about 40% of the difference between the saturated ('stable lights') and unsaturated ('radiance-calibrated data') at the national level. Note that in many ways, this understates the severity of the correction in other countries or regions. Germany is a decentralized territorial state. Many other countries in Europe are more densely populated (e.g. the Netherlands) or more centralized (e.g. France) which drives up top-coding in some areas and hence also the extent of our correction.

Next we provide 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.

Finally, we turn to three important economic applications to study the impact of top-coding and inter-calibration. The first application revisits the seminal paper by Henderson et al. (2012) and shows that the radiance-calibrated data works better for OCED countries, while current calibration methods remove the underlying economic trend alongside the "noise". The second application turns to regional inequalities as in Alesina et al. (2016) [tbc] and the third application evaluates the economic significance

of urban regions and capital cities as in Storeygard (2016) [tbc].

The paper proceeds as follows. Section 2 illustrates the extent of top-coding at the cell level. Section 3 establishes that the upper tail of the night lights is in fact Pareto distributed and presents our top-coding correction. Section 4 analyzes between and within satellite measurement errors. In Section 5 we present preliminary results of the impact of top-coding and calibration on important research questions. Section 6 concludes.

2 The Extent of Top-Coding in Worldwide Lights at the Cell Level

How important is the top coding issue? A comparison of the saturated, stable lights with the unsaturated, radiance-calibrated data can be conducted for the seven years (out of 22) for which both data exist. For illustration, we select the year 2010.[5] Our analysis covers 246 countries and territories of the earth. While light is measured at the pixel level, in this data set it is aggregated to the cell level where each cell is 0.5×0.5 decimal degrees (about 56 km^2 at the equator). Tiny states such as Gibraltar consist of one cell, while at the other extreme Russia comprises 12,324 cells. In total, we have 70,894 cells around the world, for which we have saturated and radiance calibrated light data.[6]

Table 1: Stable lights v radiance-calibrated data, worldwide, 2010, cell level

		All Cells		*Only Cells With Nonzero Lights*	
		Saturated	Rad.-Cal.	Saturated	Rad.-Cal.
Average Light Intensity in Cell	World Mean	2.3380	2.5021	4.5808	4.9046
	World St.Dev.	6.2078	11.5521	8.0766	15.8054
	World Min.	0.0000	0.0000	0.0008	0.0008
	% Cells at Min.	*48.96*	*48.98*	*0.00*	*0.00*
	World Max.	63.0000	649.9613	63.0000	649.9613
	% Cells at Max.	*0.00*	*0.00*	*0.00*	*0.00*
Maximum Light Intensity in Cell	World Mean	16.8994	30.3922	33.1112	59.5743
	World St.Dev.	22.6879	87.7104	21.7198	115.5057
	World Min.	0.0000	0.0000	3.0000	2.8196
	% Cells at Min.	*48.96*	*48.98*	*0.00*	*0.00*
	World Max.	63.0000	2415.6758	63.0000	2415.6758
	% Cells at Max.	*5.39*	*0.0000*	*10.56*	*0.0000*
# Cells		70894	70894	36183	36167
Rad.-Cal. Max. of Cells with Saturated Max. Light of 63 (World Mean)					186.3893

[5]The saturated data are averaged across the whole year, while the radiance calibrated data come from satellite F16_20100111-20101209_rad_v4, which circled the earth from 11 January to 9 December, 2010.

[6]The number of light pixels in each cell varies between 1 and 3600, with a mean of 2959.39.

Table 1 shows world summary statistics of the average and maximum light intensities in each cell in 2010, comparing the saturated 'stable lights' data (first column) to the radiance-calibrated, unsaturated numbers (second column). While the world mean of the saturated and radiance-calibrated lights are rather close (2.34 vs 2.50 for average light intensity in cell), the standard deviation of the latter is twice as high. It does not suffer from top-coding and is measured on a much wider range. By construction, the world maximum of the saturated average light intensity lies at 63 DN, the maximum of the scale, but the corresponding radiance-calibrated value is about ten times as high (649.96). The contrast becomes even starker, when we examine the maximum rather than average light intensities measured in each cell (lines 7-12). As most cells contain both bright and dim spots, the average confounds both and is not the most suitable indicator for the range and values of top lights. The world maximum of the radiance calibrated maximum light intensities of all cells is 2415.68 – 80 times as high the top coding threshold of the saturated data! Even if an observation at this magnitude might be dismissed as an outlier,[7] there is a total of 5.39% of all cells worldwide with a maximum light intensity at the top-coding threshold of 63. In fact, even this number is an understatement. Around the world, there is an enormous number of cells without any light at all; 48.96% of all cells, or 34711, have both an average and maximum light intensity of zero. In an analysis of global growth, distribution and development these (mostly) uninhabited wide plains without the faintest trace of light have to be neglected. Focusing on the remaining 36167 cells with at least one non-zero light in columns 3 and 4, we see that (i) the world means of average and maximum light intensities are doubled, with radiance-calibrated maximum light intensity showing a world mean of 59.57 and standard deviation of 115.51, (ii) the percentage of cells affected by top-coding at 63 also doubles to 10.56%. Hence, when looking only at cells with nonzero light, every tenth cell contains at least one pixel which has or exceeds a luminosity of 63 DN. Calculating the radiance-calibrated maximum light intensity of these cells gives a global mean of 186.39, three times the threshold. So both the number of cells affected and the extent of top coding are considerable.

Which countries are most affected by top-coding? Table 2 presents the figures for seven selected countries, while results for all 265 countries and territories individually are contained in Table A-1 in the Appendix. It becomes clear that countries which are predominantly (i) small (ii) highly developed and (iii) have a large degree of urbanization have particularly many cells where the maximum saturated luminosity is at the 63 DN threshold. In Israel and Belgium, this applies to 48% and 44% of all nonzero cells respectively. In these densely populated countries all cells tend to be lit, while in vast countries with areas of wilderness such as USA, Brazil and China, cells with at least one pixel larger than zero represent only between 60% and 70% of the total number of country

[7] The cell of interest lies in Saudi-Arabia, pointing to the possibility of a gas flare at the origin of the brightness.

Table 2: Stable lights v radiance-calibrated data, selected countries, 2010, cell level

	USA	Brazil	Israel	Iran	Nigeria	China	Belgium
Number of Nonzero Cells	3603	1962	23	641	446	2782	27
Nonzero Cells in %	71.26	64.41	100.00	90.69	82.60	67.62	100.00
Saturated							
Max. LI (Country Mean)	43.72	36.57	48.00	40.75	23.16	36.98	57.44
Max. LI (Country Max)	63.00	63.00	63.00	63.00	63.00	63.00	63.00
Radiance Calibrated							
Max. LI (Country Mean)	98.73	50.96	246.93	83.32	23.52	75.55	145.01
Max. LI (Country Max)	1904.87	648.63	1099.83	1253.87	459.58	1926.59	465.11
Average Rad-Cal. Max. LI if Saturated Max. LI of 63	184.07	230.27	464.16	140.46	157.99	52.00	236.31
% of Cells with Saturated Max. Light Intensity of 63	*21.59*	*10.45*	*47.83*	*20.28*	*3.93*	*9.06*	*44.44*

cells. Still, also in these countries we find a non-negligible number of cells affected by top-coding. In the U.S., nearly 22% of nonzero cells have a maximum at the 63 DN threshold. The radiance calibrated data report their mean at 184, with the maximum reaching up to 1905. The brightest pixel in China is even slightly brighter. The figures for Iran and Nigeria further underline that top-coding is not only an issue for rich countries. Sensor saturation occurs at such a low level that virtually all the countries of the world have at least one cell where some pixels are affected by it.

3 Correcting for Top-Coding

Given the importance of top-coding for the analysis of spatial inequality and development based on lights, how can we correct for it in a simple yet appropriate fashion? We suggest a top-coding correction for the whole panel, making use of the years for which we have radiance-calibrated data and applying methods from the top income literature. The histogram in Figure 3 of the distribution of the radiance-calibrated maximum light intensities of the cells affected by top-coding demonstrates the analogy. As with top incomes, the range is broad but only very few top-coded cells actually have a maximum light intensity of around 2000; the majority is concentrated in the region of 100 to 300.[8]

The procedure we outline below in effect joins the bottom and the top part from two different data sources to obtain a top-coding corrected Gini coefficient of lights and average light intensities. It thus mirrors the merging of income surveys and tax records in the top incomes literature. In this way we can easily calculate light Ginis for all the countries of the world. We illustrate our procedure with pixel level data for Germany. Germany is a highly developed country with a fairly high but not above-average rate of

[8]Note that there are 1251 out of top-coded 3821 cells which have a radiance-calibrated value below 63, which points to the difficulty of a one-to-one comparison due to the noise involved.

urbanization. The worldwide dataset shows that 25% of all nonzero cells in Germany in 2010 contained at least one pixel at the 63 DN threshold.

Figure 3: Distribution of Radiance-Calibrated Maximum Light Intensities of Cells Affected by Top-Coding (Saturated Maximum Light Intensity = 63 DN)

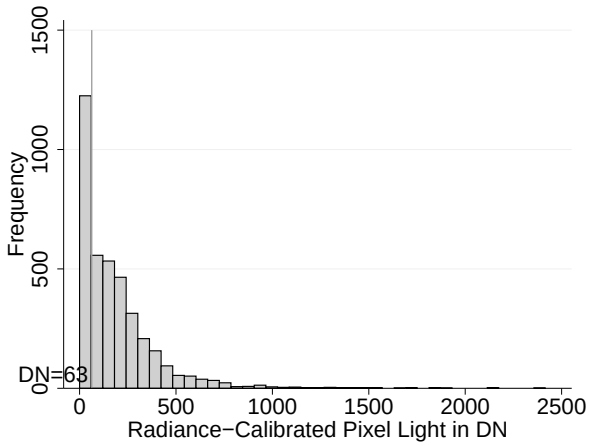

3.1 Top-Coding in Germany

We begin by presenting several summary statistics of the 'stable lights' series for Germany at the pixel level over the period from 1992 to 2013. Table 3 yields several insights: (i) The mean luminosities of German nonzero pixels fluctuate between 11 and 17 DN, with no trend discernible but quite some variation across years and satellites. (ii) The number of individual pixels affected by top-coding are small and below 1%. However, due to their geographical dispersion, this translates into 25% of cells in Germany containing at least one top-coded pixel, as discussed earlier. (iii) Inequality in lights, as measured by the Gini coefficient of these saturated data, is relatively stable across the years, fluctuating between 0.38 and 0.44.

What difference does it make when we now look at radiance-calibrated data for Germany as in Table 4? With only seven satellite observation periods, which also do not always align with calendar years, we have to be cautious with a one-to-one comparison to the saturated light series. Still, the summary statistics reveal some striking results regarding spatial inequalities. *The Gini coefficients based on radiance calibrated data are around 10 percentage points higher than those suggested by the stable lights data.* Thus, although only a small number of pixels are subject to top-coding, their influence on inequality seems to be significant and they drive up mean luminosities by a considerable amount. Hence, the saturated lights severely underestimate both mean light intensity and inequality in lights.

What would be the first idea that one might have to deal with top-coding in the years where both the saturated and radiance-calibrated lights are observed? Simply replacing

Table 3: Summary Statistics of the Saturated Data for Germany, Pixel Level

Satellite	Mean (if $x > 0$)	% with DN=63	Gini	Satellite	Mean (if $x > 0$)	% with DN=63	Gini
F101992	13.74	0.29	0.3848	F152002	14.34	0.17	0.3885
F101993	11.93	0.15	0.4038	F152003	10.57	0.04	0.4218
F101994	11.91	0.13	0.3964	F152004	10.01	0.05	0.4352
F121994	16.02	0.40	0.3607	F152005	11.75	0.11	0.4293
F121995	14.25	0.34	0.4040	F152006	11.11	0.06	0.4499
F121996	14.41	0.21	0.3913	F152007	10.79	0.09	0.4361
F121997	14.34	0.19	0.3902	F152008	17.53	0.60	0.3640
F121998	16.30	0.22	0.3800	F162004	11.91	0.12	0.3944
F121999	17.49	0.62	0.3716	F162005	10.94	0.13	0.4318
F141997	12.20	0.25	0.4173	F162006	12.37	0.07	0.4072
F141998	13.48	0.07	0.4051	F162007	12.76	0.16	0.4080
F141999	12.68	0.10	0.4088	F162008	13.07	0.17	0.4204
F142000	13.21	0.27	0.4146	F162009	13.77	0.27	0.3966
F142001	12.85	0.31	0.4219	F182010	23.34	0.65	0.3285
F142002	14.47	0.22	0.4103	F182011	14.62	0.26	0.4050
F142003	13.37	0.32	0.4314	F182012	17.42	0.54	0.3778
F152000	14.05	0.03	0.3849	F182013	14.97	0.23	0.3875
F152001	14.53	0.26	0.3974				

Table 4: Summary Statistics of the Radiance Calibrated Data for Germany, Pixel Level

Satellite Obs. Period	F12 (96/97) 16 Mar 96 - 12 Feb 97	F12 (99) 19 Jan 99 - 11 Dec 99	F12 (00) 03 Jan 00 - 29 Dec 00	F14 (02/03) 30 Dec 02 - 11 Nov 03	F14 (04) 18 Jan 04 - 16 Dec 04	F16 (05/06) 28 Nov 05 - 24 Dec 06	F16 (10) 11 Jan 10 - 9 Dec 10
Mean (if $x > 0$)	17.69	18.43	19.09	22.01	19.49	18.93	20.62
St. Dev.	26.64	31.40	33.81	36.48	33.21	28.08	25.91
Minimum	3.29	0.23	0.55	4.52	3.91	3.39	3.83
Maximum	435.25	1129.49	825.16	965.85	1045.85	500.00	487.51
99 Percentile	139.17	153.72	184.11	189.89	174.47	150.58	144.44
Gini	0.5059	0.4971	0.5073	0.5191	0.5276	0.5045	0.4577
Theil	0.5441	0.5763	0.6155	0.6038	0.6274	0.5367	0.4260
# Nonzero Obs	560175	579375	595639	578815	606840	556996	602719

the pixels with at the top-coding threshold by their radiance-calibrated counterparts. However, working with the raw values of the radiance calibrated lights is practical for two reasons, (i) the instability of the values across years and satellites, (ii) the different data ranges between saturated and radiance calibrated lights. For (i), consider Table 4. The maximum pixel luminosity in Germany (as measured by satellite F14) in 2004 is 1045 but it is only 500 a mere one year later in 2005/06 (as measured by satellite F16). 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.[9] To illustrate argument (ii), we regress the saturated lights on the radiance calibrated one for the years where both are available.[10] Table 5 We restrict the data range to the nonzero pixels with a luminosity smaller than 60, i.e. those thought to be mostly unaffected by top coding.

[9]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 Section 4 of this paper.

[10]As Table 4 shows, the radiance calibrated satellites have slightly different observation periods than the calendar years of the saturated lights. For our comparative analysis we work with the calendar years in which the vast majority of radiance calibrated data fall, e.g. year 2006 for Satellite F16 (05/06) (28 Nov 2005 - 24 Dec 2006).

Hence, an equivalence between saturated and radiance calibrated data should be possible. Instead, the regression coefficient is only of the magnitude of around 0.5 and the constant is significantly positive. The radiance-calibrated data are overall more spread out and the average (non-top-coded) pixel has a radiance calibrated luminosity of more than twice the saturated value.[11]

Table 5: Regression of Saturated on Radiance Calibrated Data for Germany, Pixel Level

Year	1996	1999	2000	2003	2004	2006	2010
Rad. Calibrated	0.5409	0.5925	0.4397	0.4219	0.3376	0.4344	0.6531
	(0.0003)	(0.0005)	(0.0003)	(0.0003)	(0.0002)	(0.0002)	(0.0006)
Constant	5.3452	7.0599	5.2533	4.6301	3.6077	4.3855	10.8590
	(0.0075)	(0.0111)	(0.0092)	(0.0087)	(0.0055)	(0.0070)	(0.0151)
R^2	0.8562	0.7544	0.7551	0.7942	0.8696	0.8510	0.6386

We draw two conclusions from our comparison: First, not all of the difference between saturated and radiance-calibrated lights can be due to top-coding. While this complicates the direct comparison, we will see that appropriately correcting for top-coding does account for a sizable proportion of the difference and produces relatively stable results. Second, rather than using the raw values of the radiance-calibrated satellites, we circumvent their instability by working only with the shape parameters of their distributions.

3.2 Finding the Top-Coding Threshold

The question of where to put the top-coding threshold is more intricate that 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.

Consider the histogram in Figure 4 with saturated values larger than 40 DN in the year 2008. The decreasing shape of the histogram, which implies that there are fewer observations with higher values, only holds till the mid-50s: Remarkably many pixels cluster at saturated light values between 55 and 63. If only the values at exactly 63 were affected by top-coding we would expect a declining number of pixels up and including to 62 DN and only a bulk at 63. Further evidence for this argument is provided by Figure 5b: Looking just at radiance calibrated top lights (at a really high value of 200 and above) in Germany in 2010, we see that they are associated with a saturated value of not only 63 DN, but there are a considerable number of 62s, 61s and all the way down to the

[11] While this may be in part owed to slight displacements of the pixels between the satellites inducing classic attenuation bias, we find it unlikely that it induces such strong variation. Increasing the size of the pixels by some scale factor, say 5, alleviates this problem somewhat but also averages out all of the interesting data points at the top.

Figure 4: Satellite F162008 with Histogram (values larger than 40) and Top 1 and 2% Correction, Germany, Pixel Level

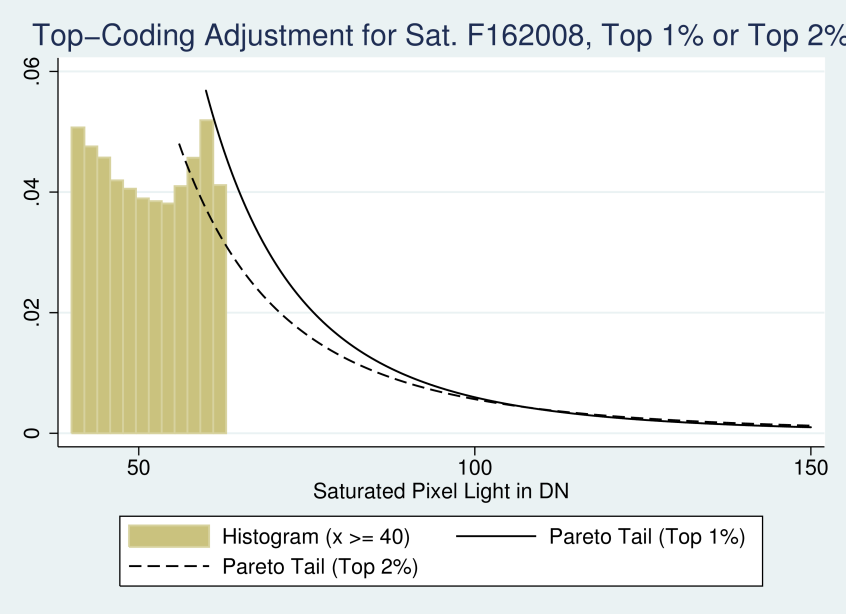

mid-50s. Figure 5a reverses the focus and depicts the distribution of the corresponding radiance-calibrated lights for all those pixels with a saturated value of 55 DN or above. While some radiance-calibrated values are very low, we do note the typical shape of the Pareto distribution among the high values.

Figure 5: The Upper Part of the Light Distribution in Germany in 2010

(a) Rad. Cal. Lights if Stable Lights $\geq$ 55 DN

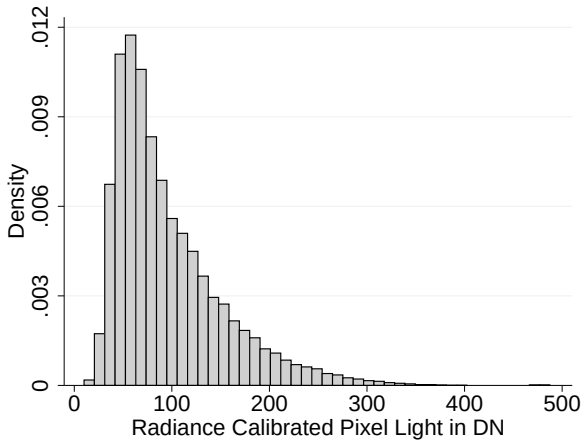

(b) Stable Lights if Rad. Cal. Lights $\geq$ 200 DN

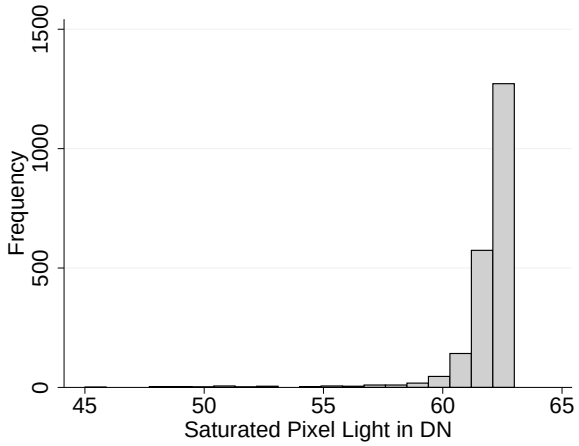

While we will examine the Pareto property of top lights in the next subsection, let us conclude from this analysis that the top-coding threshold should be set below 63 DN. Still, in percentage terms of overall nonzero pixels, we talk about the very top of the distribution: The number of nonzero pixels with the exact value of 63 DN is below 1% and including lower values down to the mid-50s can be thought to raise the percentage

of top-coded pixels to 2%. For Germany, these are around 12,000 pixels.

3.3 Are Top Lights Really Pareto Distributed?

We now present our top-coding correction procedure based on the Pareto distribution: We use the radiance calibrated data to estimate the Pareto parameter, based on which we can then extrapolate the saturated data beyond the top-coding threshold.

Of course, the first thing to do is to confirm our hypothesis and establish that there is a Pareto distribution of top lights. While a Pareto tail in the light distribution has not yet been examined, it is a standard feature of the top income literature, 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.

In fact Vilfredo Pareto (1897) first discovered that top incomes above a threshold y_c tend to follow the cumulative distribution (CDF)

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

and probability density function (PDF)

$$f(y) = \alpha \cdot y_0^{\alpha} \cdot y^{-\alpha-1} \quad \text{for} \quad y \geq y_c \tag{2}$$

with parameter $\alpha > 0$. The complement of the CDF (1) is the survival function, hence the probability that the random variable Y is larger than the given value y:

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

When taking logarithms of equation (3) – or plotting top incomes on a log scale, as Pareto did – one ends up with the characteristic linear relation:

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

This is the basis of the popular **Zipf plot**: 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, the data is concluded to be Pareto distributed. Figure 6 shows the Zipf plots of the top 2% light pixels in Germany for each of the 7 radiance calibrated satellites.[12] Despite some outliers at the very end, the Zipf plots for top lights in Germany look fairly

[12] When using not the top 2% but top 1% or top 0.5%, the plots look very similar. The same holds for the Mean Excess plot as well as the Discriminant Moment Ratio Plot, which we also repeated for alternative top percentiles.

linear and are hence indicative of the Pareto assumption.

Figure 6: Zipf Plot for Germany, Pixel Level Data, Top 2%

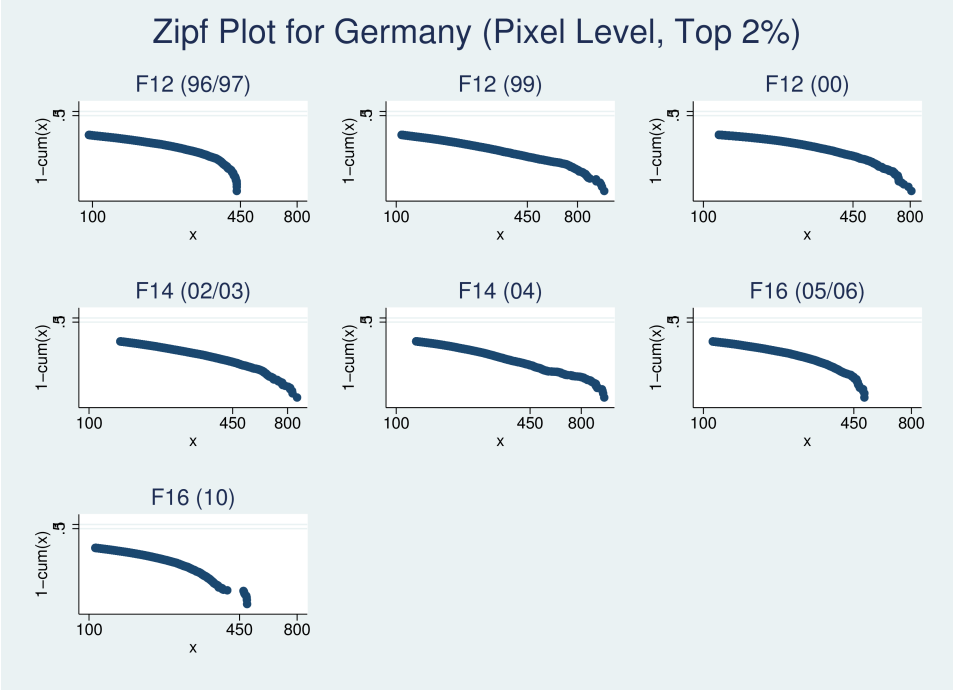

Another unique feature of the Pareto distribution is the so-called 'Van der Wijk's Law': The average income of all incomes above a given income y is proportional to y, with the factor of proportionality equal to $\frac{\alpha}{\alpha-1}$:

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

Van der Wijk's Law can be tested with a Mean Excess plot, which is shown in Figure 7 for the German top 2% light data: For each luminosity value (in the top 2%) on the x-axis, the average luminosity of all pixels brighter than this one is plotted on the y-axis. The resulting graphs look remarkably linear, with their slope hence equal to the constant factor of proportionality. With 'Van der Wijk's Law' apparently fulfilled, the Mean Excess plots provide further evidence of the Pareto distribution in top lights.

Even if Zipf plots and Mean Excess plots are standard methods for determining the Paretian nature of a data set, there has been a longstanding debate regarding their appropriateness (see already Lorenz, 1905). Cirillo (2013) demonstrates that data generated by other distributions such as the lognormal can look very similar to Paretian data in a Zipf or Mean Excess plot. She therefore proposes an additional test which does not rely on pure visual evaluation of linearity: The discriminant moment ratio plot shows the coordinate pair of coefficient of variation (this is, standard deviation divided by mean) on the x-axis and skewness on the y-axis. As each parametric distribution has

Figure 7: Mean Excess Plot for Germany, Pixel Level Data, Top 2%

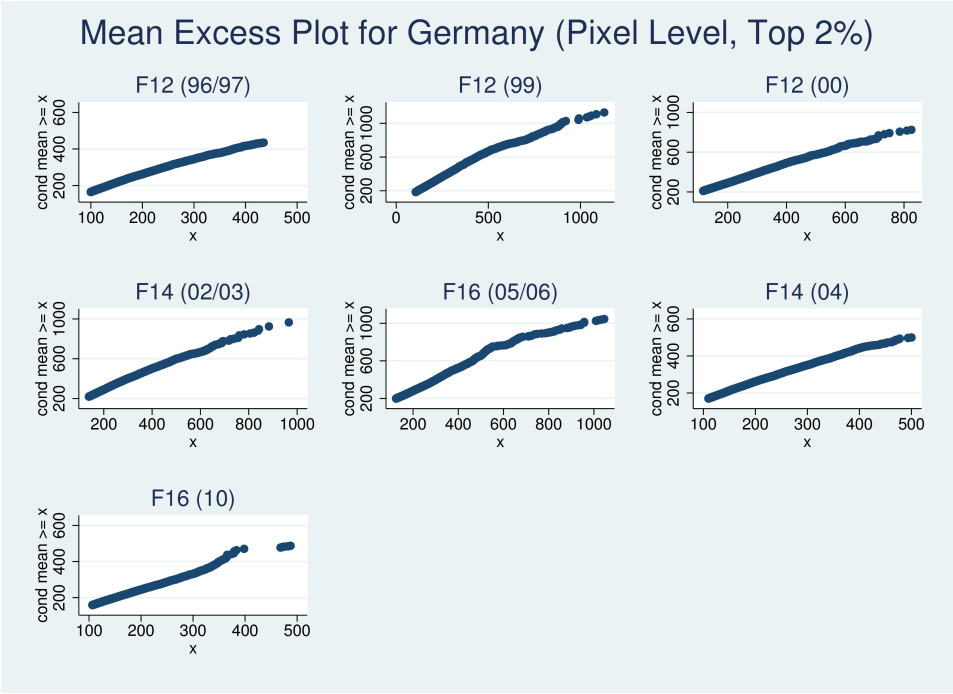

its particular curve of coordinate pairs, one can divide the area, among others, into a Paretian area, a Lognormal area and a Gray area in between. Figure 8 for Germany shows that 5 out of the 7 radiance calibrated satellites are well inside the Pareto area and the other two in the gray zone.

Hence, the Discriminant Moment Ratio Plot confirms the results from the Zipf and Mean Excess plot and we are quite safe to assume a Pareto distribution for top lights in Germany.

3.4 Augmenting the Saturated Data with a Pareto Tail

Our top-coding correction follows a three-step procedure to derive national or sub-national estimates of inequality using the lights data

1. Estimate the shape parameter α of the Top 2% using the radiance-calibrated data
2. Transfer these α estimates to the corresponding years of the saturated data, using parameter interpolation for intermediate years
3. The saturated data for all values below the threshold (98%) are combined with a Pareto tail for the top 2% based on the determined α

The maximum likelihood estimates of the Pareto parameter α of the radiance calibrated data for Germany are shown in Table 6. When assuming a Pareto distribution for the top 2%, most satellites' α parameters are between 2.2 and 2.6. For comparison,

Figure 8: Discriminant Moment Ratio Plot for Germany, Pixel Level Data, Top 2%

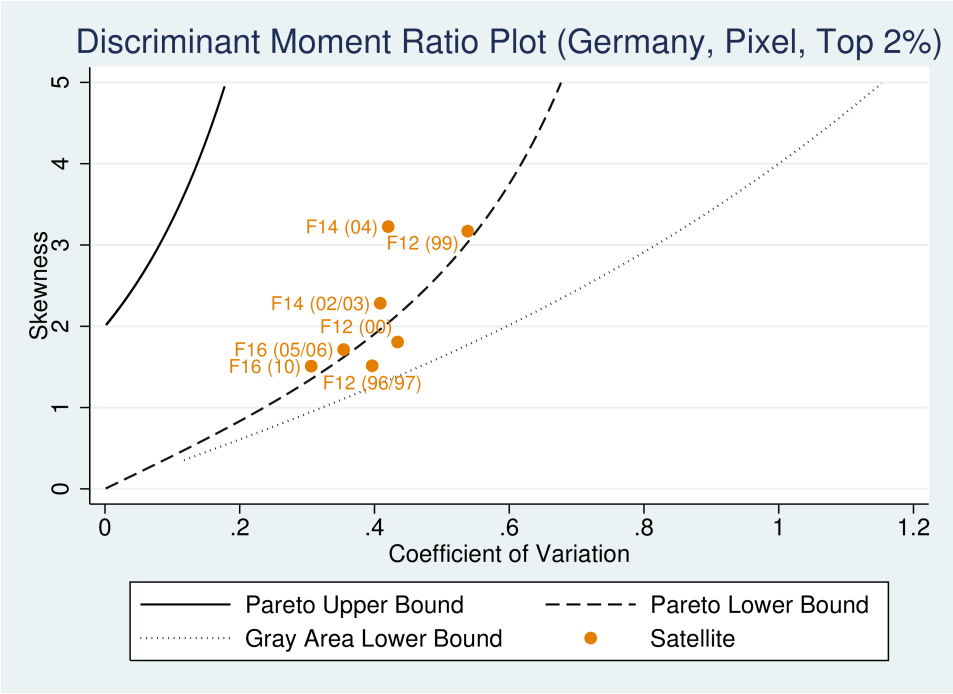

Table 6: Estimation Results of the Pareto Parameter Alpha for Different Top Percentage Levels, Germany, Pixel Level

Satellite	F12 (96/97)	F12 (99)	F12 (00)	F14 (02/03)	F14 (04)	F16 (05/06)	F16 (10)
Pareto Alpha (Top 2%)	2.2612	2.1594	1.9442	2.5176	2.4761	2.5944	2.7843
(Standard Error)	(0.0214)	(0.0201)	(0.0178)	(0.0234)	(0.0225)	(0.0245)	(0.0254)
Mean (Top 2%)	172.98	198.42	241.00	230.38	210.28	178.99	166.41
# Obs (Top 2%)	11204	11588	11914	11577	12154	11223	12055
Top 2% Threshold	96.48	106.53	117.04	138.87	125.36	110.00	106.64
Pareto Alpha (Top 1%)	2.8316	2.5309	2.7705	2.9757	3.1118	3.1849	3.6340
(Standard Error)	(0.0378)	(0.0333)	(0.0359)	(0.0391)	(0.0399)	(0.0427)	(0.0468)
Mean (Top 1%)	215.15	254.13	288.09	286.00	257.09	219.49	199.28
# Obs (Top 1%)	5602	5793	5960	5789	6068	5570	6027
Top 1% Threshold	139.17	153.72	184.11	189.89	174.47	150.58	144.44
Pareto Alpha (Top 0.5%)	3.5802	2.8452	3.5367	3.3156	3.6194	3.9211	4.6784
(Standard Error)	(0.0676)	(0.0529)	(0.0648)	(0.0616)	(0.0657)	(0.0743)	(0.0852)
Mean (Top 0.5%)	258.87	320.01	345.78	350.48	310.33	261.01	231.18
# Obs (Top 0.5%)	2801	2897	2978	2894	3035	2784	3013
Top 0.5% Threshold	186.56	207.54	248.01	244.77	224.59	194.44	181.77

the parameter estimates for the smaller top shares (1% and 0.5%) are higher and hence associated with both more probability mass near the cut-off threshold and a shorter tail. They also exhibit much larger standard errors and are based on fewer observations.

Figure 4 shows the difference between topping up the saturated data of satellite F162008 with a Pareto tail for the top 1% and top 2%.[13] As discussed above, we have

[13]Note that the matching density values of the saturated data (the histogram) and the Pareto distributions have been chosen for visual presentation. A histogram based on all values rather than just those above 40 DN would lead to lower density values but what matters is the shape. The combined semi-parametric CDF can be calculated with the formula by Dupuis and Victoria-Feser (2006):

$$F(y) = \begin{cases} G(y) \text{ if } y \leq y_c \\ G(y_c) + (1 - G(y_c)) \cdot F_\alpha(y) \text{ if } y > y_c \end{cases} \tag{6}$$

good reasons to assume a top-coding threshold rather in the 50s than at exactly 63 DN, so we henceforth conduct our top-coding procedure for the top 2% (while presenting results for the top 1% in the Appendix).

After augmenting the bottom 98% of the saturated data with a top 2% Pareto tail for each of the saturated satellite-years we arrive at our top-coding corrected Gini coefficients. One can think of the top 2% being replaced by values drawn from a Pareto distribution with the estimated shape parameter (as is common in the top incomes literature, e.g. see Alfons et al., 2013), however, in order to compute the corrected Gini coefficients we do not even have to got that far. Thanks to a simple formula for the Gini coefficient of two non-overlapping subgroups, we do not have to actually conduct the replacement of the top tail at the pixel level. The Appendix shows that the overall Gini consists of 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) (e.g. see Cowell, 2013):

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

where ω_B and ω_B are the pixel shares for the data below the threshold, denoted B, and the data above the threshold, denoted T, with $\omega_T = 1 - \omega_T$. The shares of all income (light) accruing to either data source are $\phi_B = \omega_B \mu_B / \mu$ and $\phi_T = \omega_T \mu_T / \mu$. The overall mean is simply a weighted-average of the data below the threshold and the Pareto distributed data above:

$$\mu = \omega_B \mu_B + \omega_T \frac{\alpha}{\alpha - 1} y_c \tag{8}$$

where μ_B is estimated by the sample mean below the threshold y_c.

The formulas highlight two important insights about the dynamics of the top-coding corrected Gini coefficient and mean. On the one hand, a greater share of top-coding, brighter top-coded pixels, and a greater spread in the distribution of the top-coded data all increase estimates of inequality. Unsurprisingly, the bias induced by top-coding into assessments of inequality will be larger in urbanized or highly developed areas compared to rural or less-developed regions. On the other hand, *top-coding substantially drives up differences in mean lights* in areas affected by it. For illustration, consider this simple numerical example. If top-coding affects only 5% of the study area of interest, the shape parameter is $\alpha = 2$, average light intensity in the non-top-coded pixels is 20 DN and top-coding occurs at $y_c = 63$ DN, then the corrected mean is 25.3 DN. While this example is more in line with a densely populated nation, if the analyst wishes to analyze smaller regions or study urban development, then both the share of top-coded pixels and estimates

of the shape parameter are likely to rise considerably.

Plugging a top-coded percentage of the top 2%, i.e. $\omega_T = 2\%$, into (7), it becomes obvious that the formula for the top-coding corrected Gini just depends on α, y_c, μ_B and G_B, all of which are available without replacement at the pixel level:

$$G = \frac{0.98^2 \cdot \mu_B \cdot G_B + 0.02^2 \cdot \frac{\alpha}{\alpha-1} y_c \cdot \frac{1}{2\alpha-1} + 0.02 \cdot \frac{\alpha}{\alpha-1}}{0.98 \cdot \mu_T + 0.02 \cdot \frac{\alpha}{\alpha-1} y_c} - 0.02. \tag{9}$$

Table 7: Light Gini with a Pareto Tail for the Top 2%, Germany, Pixel Level

Satellite	Gini unadjusted	α	y_c	μ_P	G_P	μ_R	G_R	G	Δ
F101992	0.3848	2.26	55	12.82	0.3559	98.61	0.2839	0.4180	0.0331
F101993	0.4038	2.26	53	10.98	0.3686	95.02	0.2839	0.4380	0.0341
F101994	0.3964	2.26	52	10.97	0.3608	93.23	0.2839	0.4300	0.0336
F121994	0.3607	2.26	57	15.12	0.3371	102.19	0.2839	0.3922	0.0315
F121995	0.4040	2.26	57	13.31	0.3783	102.19	0.2839	0.4367	0.0327
F121996	0.3913	2.26	55	13.49	0.3656	98.61	0.2839	0.4224	0.0311
F121997	0.3902	2.23	56	13.42	0.3642	101.63	0.2895	0.4238	0.0336
F121998	0.3800	2.19	57	15.41	0.3594	104.77	0.2953	0.4119	0.0319
F121999	0.3716	2.16	60	16.59	0.3525	111.75	0.3013	0.4053	0.0337
F141997	0.4173	2.23	54	11.24	0.3847	98.00	0.2895	0.4520	0.0347
F141998	0.4051	2.19	54	12.57	0.3785	99.25	0.2953	0.4391	0.0340
F141999	0.4088	2.16	54	11.75	0.3780	100.58	0.3013	0.4450	0.0362
F142000	0.4146	1.94	56	12.26	0.3860	115.31	0.3462	0.4595	0.0449
F142001	0.4219	2.14	56	11.89	0.3924	105.32	0.3057	0.4597	0.0377
F142002	0.4103	2.33	55	13.56	0.3874	96.46	0.2737	0.4390	0.0287
F142003	0.4314	2.52	58	12.41	0.4048	96.22	0.2478	0.4598	0.0284
F152000	0.3849	1.94	55	13.14	0.3579	113.25	0.3462	0.4289	0.0440
F152001	0.3974	2.14	57	13.59	0.3718	107.21	0.3057	0.4333	0.0360
F152002	0.3885	2.33	56	13.42	0.3623	98.22	0.2737	0.4195	0.0310
F152003	0.4218	2.52	48	9.67	0.3851	79.63	0.2478	0.4477	0.0259
F152004	0.4352	2.48	49	9.09	0.3955	82.19	0.2530	0.4638	0.0285
F152005	0.4293	2.54	51	10.83	0.3981	84.22	0.2457	0.4544	0.0251
F152006	0.4499	2.59	51	10.18	0.4177	82.99	0.2387	0.4743	0.0244
F152007	0.4361	2.64	51	9.85	0.3997	82.06	0.2334	0.4608	0.0247
F152008	0.3640	2.69	59	16.63	0.3444	93.93	0.2284	0.3864	0.0224
F162004	0.3944	2.48	52	10.98	0.3590	87.23	0.2530	0.4230	0.0286
F162005	0.4318	2.54	51	10.00	0.3958	84.22	0.2457	0.4584	0.0266
F162006	0.4072	2.59	52	11.45	0.3763	84.61	0.2387	0.4321	0.0249
F162007	0.4080	2.64	55	11.81	0.3769	88.50	0.2334	0.4336	0.0256
F162008	0.4204	2.69	56	12.12	0.3920	89.15	0.2284	0.4451	0.0248
F162009	0.3966	2.74	55	12.84	0.3690	86.67	0.2235	0.4194	0.0228
F182010	0.3285	2.78	61	22.55	0.3182	95.19	0.2189	0.3468	0.0183
F182011	0.4050	2.78	57	13.69	0.3809	88.95	0.2189	0.4272	0.0222
F182012	0.3778	2.78	60	16.52	0.3591	93.63	0.2189	0.3995	0.0217
F182013	0.3875	2.78	57	14.05	0.3630	88.95	0.2189	0.4099	0.0224

Table 7 shows the Gini coefficient of the saturated data for Germany (Col.2), all

the ingredients for top-coding correction according to formula (9) as well as the resulting corrected Gini coefficient (Col.10) and the difference to the unadjusted one (last column): We see that correcting for top-coding with the help of a Pareto tail increases the Gini coefficient by 3 to 4 percentage points in earlier years and 2 to 3 percentage points in later years.[14] This makes a huge difference indeed. Remember that for the radiance calibrated satellites we found Gini coefficients which are 10 percentage points higher than the saturated ones (cf Table 4). *Simply adjusting the top 2% of the data accounts for up to 40% of the difference between saturated and radiance-calibrated Gini coefficients.*

Figure 9: Top-Coding Corrected Ginis for Germany (Averaged Across Satellites), Pixel Level, Top 2% Correction

In order to be able to trace the top-coding adjusted Gini coefficient from Table 7 over time, we average the results for those years where we have more than one satellite. The resulting time path in Figure 9 shows that inequality in light in Germany increased rather steadily in the 2000 decade before going down in 2010 and recovering in the last years. It is now up to further research to explain this development of the light Gini in Germany and to compare it to those of other countries. Based on our proposed top-coding correction procedure it is straightforward to follow the illustration of Germany and to calculate inequality in light measures for all the countries of the world from 1992 to 2013. With such a database, one can literally shine a new light on regional inequality and convergence in the world.

As we have shown above, our formula for the top-coding adjusted Gini coefficient (9)

[14]Tables Table A-4 and Figure A-1 in the Appendix show that with a Top 1% Pareto tail, the Gini correction is less pronounced but follows the same direction.

only depends on parameters from the bottom part of the saturated data and the top of the radiance calibrated ones, making it unnecessary to actually replace the values in the top tail at the pixel level. Still, there are some important questions in development and regional economics for which an overall Gini in lights is not a sufficient measure of analysis. For instance, spatial growth regressions or the identification of growth clubs require light data at the pixel level of the corresponding geographical location. Our top-coding correction procedure readily extends to that: The bottom part of the saturated data (e.g. 98%) is kept and the top 2% is actually replaced at the pixel level by sorted values drawn from a Pareto distribution with the estimated parameter, see also e.g. Alfons et al. (2013). While such a pixel-level replacement is obviously more data-intensive, it is straightforward to carry out.

4 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 10 plots the average light intensity in Sicily as obtained from each separate satellite and shows how severe these jumps can be.

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. While there is little that we can do about bottom-coding (the data are simply not observed), the aim of this paper is to provide a thorough treatment and offer solutions of how to deal

with the other two types of errors. We already addressed top-coding and now exclusively focus on shifts in brightness among the different images.

Figure 10: Average lights in Sicily according to each satellite

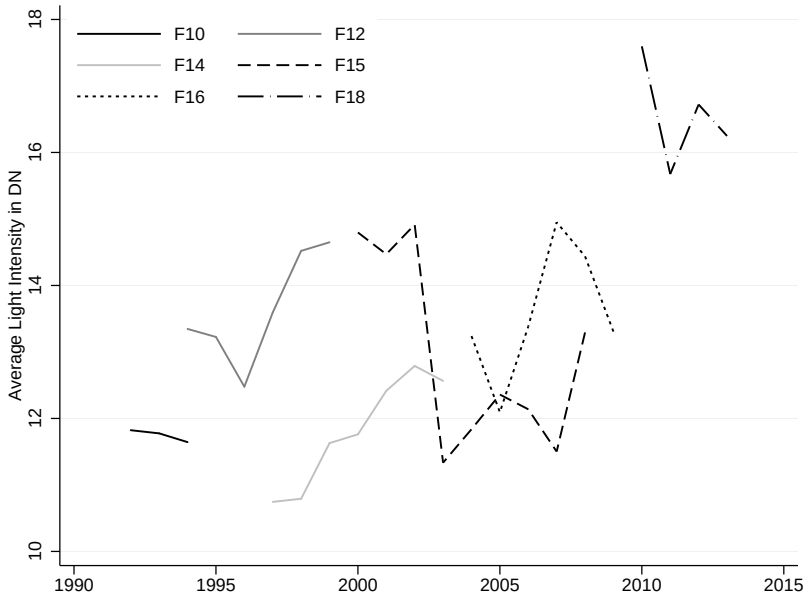

Economists usually deal with these measurement errors by including time fixed effects in the regression of interest (e.g. see Henderson et al., 2012; Michalopoulos and Papaioannou, 2013; Hodler and Raschky, 2014). Chen and Nordhaus (2011) and Henderson et al. (2012) provide a detailed discussion and 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).

We highlight the properties of the calibration in this section. The next section shows how this changes the results in several interesting applications.

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.[15] 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 in the Appendix 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,[16] 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 11: Results of "inter-satellite calibration" and real GDP in Sicily

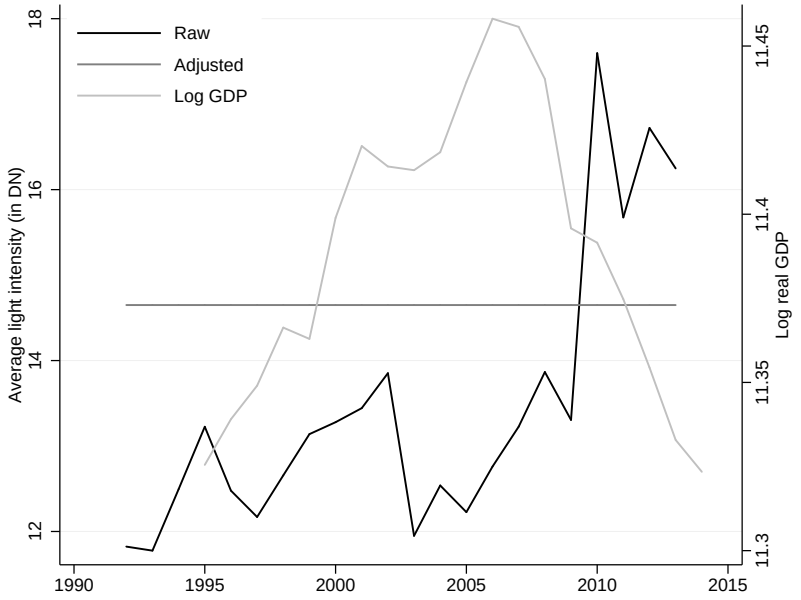

Figure 11 shows why this approach is both attractive and very problematic at the same time. It plots the average light intensity in Sicily before and after the adjustment. We simply average the data whenever we have more than one satellite at our disposal. One the one hand, the adjusted average light 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

[15]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.

[16]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.

$N \cdot \mathrm{E}[Y] = N \cdot \mathrm{E}[Y|\bar{X}]$). 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 11 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).

Even though it has empirical shortcomings, the approach proposed by Elvidge et al. (2009) is perfectly suited to adjust images to match the average and total light intensity of a particular reference satellite. Sicily is just not a good reference area. If light emissions in Sicily in part mirror economic growth, then applying this adjustment to the entire world will not only remove noise but instead remove parts of the underlying economic trend as well. Such as scenario is entire plausible. For example, an increase in average light intensity from 12 DN to 14 DN is approximately in line with a 15% increase in GDP from 1995 to 2007. We return to this question in the next section.

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 still searching for suitable candidates at the moment (but from an economist's point of view, parts of Japan may be an intriguing choice for the period starting in the 1990s). A second 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 such an application and evaluate the utility of non-linear models that respect the bound of the observed data.

5 Applications

We illustrate the economic significance of top-coding and satellite "inter-calibration" 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. The second application studies regional and ethnic inequality inspired by a recent path-breaking paper by Alesina et al. (2016). Finally,

we explore 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.

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.

5.1 Lights and GDP growth

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 8 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 8 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 top-coding than poorer countries. Since we have not yet applied our correction approach 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 8 illustrates an interesting and novel 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

Table 8: 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
Log Lights per sq. km	0.226***		0.181***	
	(0.072)		(0.059)	
Log 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

Note(s): All columns include country and time fixed-effects. Country-clustered standard errors in parentheses. Significant at: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 9 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 does not differ systematically between rich countries and poorer countries.

Having established how top-coding affects the result in Henderson et al. (2012), we now turn to how satellite inter-calibration affects the estimated output elasticities. For now we just restrict our attention to the stable lights data (since calibrating the unsaturated series is another issue). Again, we construct a simple test of the influence of calibrating the series by just re-running the original regressions with the different data sets. Note that this process is cumbersome because it requires us to adjust all images at the pixel level first, then compute area-weighted images, and finally calculate the luminosity values for each country in each satellite-year. Since the correction sometimes scales some pixels in the new image above a DN of 63, we then replace these by their top-coded counterpart.

Table 9: 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
Log Lights per sq. km	0.242***		0.184***	
	(0.072)		(0.062)	
OECD × Log Lights per sq. km	-0.237***		-0.044	
	(0.035)		(0.058)	
Log Lights per capita		0.239***		0.180***
		(0.064)		(0.056)
OECD × Log 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

Note(s): All columns include country and time fixed-effects. Country-clustered standard errors in parentheses. Significant at: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

As before, whenever we have two satellites at our disposal for any given year, we simply average the adjusted data.

Table 10 shows the corresponding results. Column (1) uses the raw stable lights data for the entire sample from 1992 to 2013. Again, we find an elasticity that is close to the usual result of about 0.3. Column (2) introduces the series that has been calibrated using the method described in the previous section and the coefficients reported in Table A-5 in the Appendix – "Adjusted (1)". The results are not good news for this type of calibration. The elasticity is now only about half of the earlier estimate and the standard error widens, so that we can no longer reject the null of zero. Since our calibration table yields different results than those reported in Elvidge et al. (2014), we try again using their coefficients – "Adjusted (2)". Column (3) uses slightly fewer observations since we now lack coefficients for two satellites but otherwise confirms our initial finding. Interestingly, the problem is much less severe when we use per capita quantities as the second half of the table shows.

Our preliminary interpretation of this finding is that the NOAA approach to calibrating satellites indeed removes much of the underlying trend in the lights series together with the "noise". As a result, the correlation between *changes* in lights and *changes* in GDP weakens substantially (although the correlation in levels rises). Once we divide both sides of the equation by population, this problem is alleviated by the

Table 10: Income elasticity of lights, calibrations, 1992-2013, country-level

	GDP v Lights per sq. km			*GDP per capita v Light per capita*		
	(1)	(2)	(3)	(4)	(5)	(6)
Raw data	0.316***			0.298***		
	(0.078)			(0.071)		
Adjusted (1)		0.140			0.259**	
		(0.098)			(0.102)	
Adjusted (2)			0.175			0.310**
			(0.134)			(0.120)
Constant	25.812***	24.972***	24.724***	11.843***	10.235***	10.072***
	(0.024)	(0.513)	(0.741)	(0.310)	(0.109)	(0.169)
Within-R2	0.748	0.677	0.673	0.563	0.475	0.471
Observations	4198	4199	4011	4195	4196	4008
Countries	199	199	199	199	199	199

Note(s): All columns include country and time fixed-effects. Country-clustered standard errors in parentheses. Significant at: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

introduction of additional variation. This confirms an interesting trend across all these tables, namely per capita quantities deliver much more stable estimates of the underlying relationship than scale dependent quantities, such as GDP in constant prices.

5.2 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. That point should already be clear from the discussion of Germany above where the Gini differs by about 3-4 points even though it is not too densely populated. 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 base specification at the country-level.]

5.3 Urban-rural differences

[To be completed later... We plan to use two different ways of isolating the economic significance of cities 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) use 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.]

6 Conclusion

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, leading to a upward bias in estimates of regional convergence and a downward bias in estimates of spatial inequality.

We show that top lights, just as top incomes, are Pareto distributed and suggest a solution based on methods from the top income literature: We augment the lower part of the saturated data with a Pareto tail based on an estimated α parameter from the unsaturated, radiance calibrated data. 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 Germany show that top-coding correction makes a substantial difference, raising the estimated Gini coefficient of lights by 2 to 5 points and accounting for up to 40% of the difference between the saturated ('stable lights') and unsaturated ('radiance-calibrated data') at the national level.

For the next version of this paper, we plan to repeat this Gini top-coding correction for all other countries, resulting in a worldwide panel data set of top-coding corrected inequality measures from 1992-2013 to be used for further analysis.

For other applications where location specific light measures rather than a national Gini index is required, we intend to conduct the top-coding correction at the pixel level involving the replacement of pixels by those sampled from the Pareto distribution. While this is obviously more data-intensive, it is straightforward to carry out, see e.g. Alfons et al. (2013) for a similar application using income surveys.

Furthermore, in the future we also aim to provide a better approach to solve the satellite inter-comparability problem. Having demonstrated that Elvidge et al.'s (2009) suggestion of scaling the various images to a reference area satellite and reference area of Sicily is flawed, it is now up to further research to find a better way that preserves the time series dimension of the lights data.

Our first applications to determine for which question in development and urban

economics the top-coding problem is most severe also opens the door to further research. One of our most intriguing findings 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. It can be considered a riposte to studies which argue that nighttime lights are an inadequate income proxy in developed countries. 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 and open up many new ways of looking at inequality and development on a global scale.

References

Alesina, A. F., S. Michalopoulos, and E. Papaioannou (2016). Ethnic inequality. *Journal of Political Economy* (forthcoming).

Alfons, A., M. Templ, and P. Filzmoser (2013). Robust estimation of economic indicators from survey samples based on pareto tail modelling. *Journal of the Royal Statistical Society: Series C (Applied Statistics) 62*(2), 271–286.

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.

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.

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.

Dreher, A. and S. Lohmann (2015). Aid and growth at the regional level. *Oxford Review of Economic Policy 31*, 420–446.

Dupuis, D. and M.-P. Victoria-Feser (2006). A robust prediction error criterion for pareto modelling of upper tails. *Canadian Journal of Statistics 34*, 639–658.

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.

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.

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.

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 (2014). Lights, camera,... income! Estimating poverty using national accounts, survey means, and lights. *NBER Working Paper w19831*.

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.

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

A Appendix

A.1 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

$$G = \frac{1}{2N^2\mu}\sum_i\sum_j |y_i - y_j| \tag{A-1}$$

$$= \frac{1}{2N^2\mu}\sum_k\left(\sum_{i\in N_k}\sum_{j\in N_k}|y_i - y_j| + \sum_{i\in N_k}\sum_{j\notin N_k}|y_i - y_j|\right) \tag{A-2}$$

$$= \sum_k\left(\frac{N_k}{N}\right)^2\frac{\mu_k}{\mu}G_k + \frac{1}{2N^2\mu}\sum_k\sum_{i\in N_k}\sum_{j\notin N_k}|y_i - y_j|\,. \tag{A-3}$$

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 $\sum_{i\in N_k}\sum_{j\in N_h}|y_i - y_j| = N_kN_h\,|\mu_k - \mu_h|$. Hence, we can simplify equation (A-3) to

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

With two 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 $\phi_B = \omega_B\mu_B/\mu$ and $\phi_T = \omega_T\mu_T/\mu$. Last, recognize that G_P is computed using the non-censored data ($y \in [0, y_c)$) and G_B is the Gini of the top coded data ($y \in [y_c, \infty]$) which we assume to be Pareto distributed. This implies $G_T = 1/(2\alpha - 1)$ and $\mu = \omega_B\mu_B + \omega_T\alpha/(\alpha - 1)y_c$. Then we obtain equation (7) from the main text

$$G = \omega_B\phi_BG_B + \omega_T\phi_T\frac{1}{2\alpha - 1} + [\phi_T - \omega_T] \tag{A-6}$$

which can be estimated easily using only the top-coded data, the censoring cut-off y_c and the estimated Pareto coefficient α.

A.2 Additional Tables and Figures

Figure A-1: Top-Coding Corrected Ginis for Germany (Averaged Across Satellites), Pixel Level, Top 1% Correction

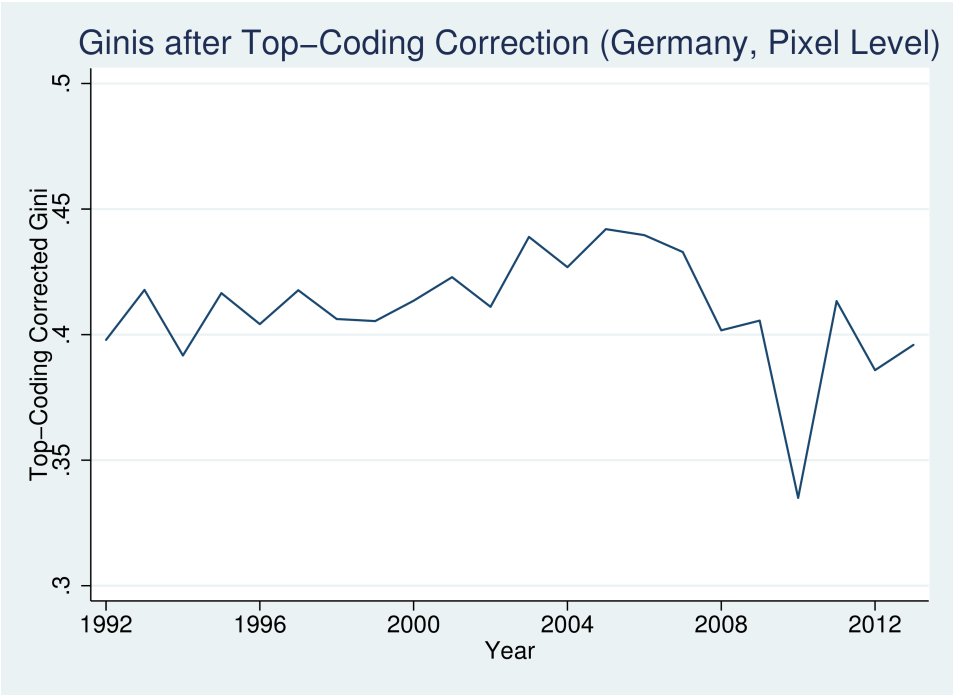

Table A-1: Number of Non-Zero Cells Affected by Top Coding in Each Country (Year 2010), Part 1

Country	# Nonzero	% of Nonzeros Top-Coded	Top-Coded Mean of Max LI	Top-Coded Max of Max LI
ALA	5	20.00	154.46	154.46
AFG	112	2.68	180.37	201.74
ALB	20	10.00	227.36	279.47
DZA	328	11.28	292.25	2132.96
ASM	4	0.00		
AND	1	100.00	223.60	223.60
AGO	174	4.02	121.86	282.21
AIA	1	0.00		
ATA	3	0.00		
ATG	3	0.00		
ARG	858	7.46	112.40	728.46
ARM	21	9.52	70.38	134.22
ABW	1	0.00		
AUS	0			
AUS	963	3.22	19.15	101.17
AUT	63	9.52	210.74	284.59
AZE	61	3.28	387.76	391.55
BHS	39	5.13	14.79	29.59
BHR	2	100.00	956.70	979.85
BRN	6	16.67	316.66	316.66
UMI	0			
BGD	74	4.05	130.82	208.28
BRB	1	100.00	145.69	145.69
BLR	142	9.86	319.40	857.25
BEL	27	44.44	236.31	465.11
BLZ	12	0.00		
BEN	30	3.33	116.22	116.22
BMU	1	0.00		
BTN	14	0.00		
BOL	183	6.01	304.82	492.90
BES	3	0.00		
BIH	35	2.86	150.78	150.78
BWA	92	1.09	161.44	161.44
BVT	1	0.00		
BRA	1962	10.45	230.27	648.63
IOT	1	0.00		
VGB	2	0.00		
BGR	67	2.99	121.01	214.05
BFA	54	5.56	113.72	135.56
MMR	225	1.78	54.84	189.15
BDI	13	0.00		
KHM	70	1.43	49.37	49.37
CMR	91	2.20	152.74	167.05
CAN	1802	6.33	161.61	862.46
CPV	12	0.00		
CYM	3	33.33	145.25	145.25
CAF	30	0.00		
TCD	63	1.59	103.44	103.44
CHL	237	10.55	70.95	277.54
CHN	2782	9.06	52.00	855.42
CXR	1	0.00		
C–	0			
CCK	0			
COL	254	8.27	269.74	516.20
COM	5	0.00		
COG	47	4.26	144.95	171.07
COD	124	2.42	156.45	192.33
COK	2	0.00		
AUS	0			
CRI	33	12.12	234.84	313.89
HRV	56	8.93	184.85	314.41
CUB	68	4.41	143.69	208.74
CUW	1	100.00	205.92	205.92
CYP	9	55.56	233.78	393.68
CZE	57	12.28	203.13	291.41
CIV	109	2.75	166.18	194.89
DNK	52	28.85	111.13	296.96
DJI	5	20.00	100.89	100.89
DMA	1	0.00		
DOM	30	10.00	267.07	419.87
ECU	77	9.09	202.16	414.18
EGY	186	29.57	209.82	940.31
SLV	13	15.38	205.36	290.85
GNQ	14	0.00		
ERI	18	0.00		
EST	47	12.77	522.36	929.22
ETH	169	0.59	134.15	134.15
ATF	0			
FLK	12	0.00		
FRO	5	0.00		
FSM	4	0.00		
FJI	18	0.00		
FIN	272	12.13	274.22	640.80
FRA	324	29.94	187.71	874.55
GUF	13	0.00		
PYF	30	0.00		
ATF	7	0.00		
GAB	46	0.00		
GMB	10	0.00		
PSE	7	85.71	294.09	679.34
GEO	41	2.44	263.22	263.22
DEU	230	24.78	138.96	487.51
GHA	81	6.17	121.13	182.31

Table A-2: Number of Non-Zero Cells Affected by Top Coding in Each Country (Year 2010), Part 2

Country	# Nonzero	% of Nonzeros Top-Coded	Top-Coded Mean of Max LI	Top-Coded Max of Max LI
GIB	1	0.00		
ATF	0			
GRC	141	4.26	290.29	636.32
GRL	22	0.00		
GRD	2	0.00		
GLP	4	25.00	181.48	181.48
GUM	1	100.00	143.42	143.42
GTM	45	2.22	365.66	365.66
GGY	1	0.00		
GIN	42	0.00		
GNB	5	0.00		
GUY	17	0.00		
HTI	17	0.00		
HMD	2	0.00		
HND	51	1.96	244.49	244.49
HKG	2	50.00	953.36	953.36
UMI	0			
HUN	63	11.11	174.10	366.56
ISL	68	1.47	62.84	62.84
IND	1200	14.67	66.03	499.33
IDN	644	3.57	59.45	632.81
IRN	641	20.28	140.46	813.13
IRQ	136	19.85	311.26	2149.87
IRL	60	10.00	304.86	620.18
IMN	1	0.00		
ISR	23	47.83	464.16	1099.83
ITA	225	24.00	242.33	636.59
JAM	9	22.22	172.15	236.12
JPN	288	25.69	376.45	1842.57
UMI	0			
JEY	1	0.00		
UMI	0			
JOR	33	18.18	321.89	658.07
ATF	0			
KAZ	897	2.01	223.24	640.06
KEN	80	2.50	141.91	164.13
UMI	0			
KIR	4	0.00		
KWT	13	38.46	634.03	1288.69
KGZ	88	1.14	384.88	384.88
LAO	71	1.41	157.03	157.03
LVA	54	7.41	224.01	454.12
LBN	9	33.33	303.22	475.53
LSO	13	7.69	77.16	77.16
LBR	20	0.00		
LBY	212	11.79	150.50	837.37
LIE	1	0.00		
LTU	51	13.73	168.91	266.04
LUX	3	33.33	213.08	213.08
MAC	1	100.00	398.17	398.17
MKD	15	13.33	205.86	366.61
MDG	71	1.41	76.56	76.56
MWI	42	2.38	92.45	92.45
MYS	147	22.45	187.10	1524.76
MDV	25	0.00		
MLI	76	2.63	149.98	150.88
MLT	1	100.00	349.44	349.44
MHL	8	0.00		
MTQ	1	100.00	263.68	263.68
MRT	42	2.38	183.64	183.64
MUS	4	50.00	133.33	143.94
MYT	1	0.00		
MEX	753	16.87	248.41	1846.80
UMI	0			
MDA	31	0.00		
MCO	1	0.00		
MNG	158	0.63	29.87	29.87
MNE	11	0.00		
MSR	1	0.00		
MAR	161	9.32	141.17	620.97
MOZ	130	0.77	214.33	214.33
NAM	142	0.70	174.25	174.25
NRU	1	0.00		
UMI	0			
NPL	42	2.38	7.90	7.90
NLD	35	42.86	234.17	437.68
NCL	22	4.55	169.02	169.02
NZL	125	4.00	9.73	44.63
NIC	47	2.13	213.74	213.74
NER	75	1.33	97.76	97.76
NGA	280	3.93	157.99	450.62
NIU	1	0.00		
NFK	1	0.00		
PRK	51	0.00		
MNP	3	0.00		
NOR	334	12.28	249.36	703.33
OMN	115	20.87	201.73	603.67
PAK	284	8.45	243.27	961.49
PLW	2	0.00		
VUT	9	0.00		
UMI	0			
PAN	36	5.56	204.00	241.65
PNG	60	1.67	111.98	111.98

Table A-3: Number of Non-Zero Cells Affected by Top Coding in Each Country (Year 2010), Part 3

Country	# Nonzero	% of Nonzeros Top-Coded	Top-Coded Mean of Max LI	Top-Coded Max of Max LI
PRY	100	6.00	269.54	455.01
PER	285	3.16	79.12	343.49
PHL	197	5.08	101.60	262.75
PCN	0			
POL	200	31.00	229.44	807.23
PRT	74	12.16	213.10	414.27
PRI	9	66.67	246.54	512.63
QAT	8	37.50	585.79	1281.62
REU	3	33.33	203.26	203.26
ROU	134	9.70	181.36	446.95
RUS	4645	3.77	16.25	383.58
RWA	10	0.00		
BLM	1	0.00		
SHN	3	0.00		
KNA	1	0.00		
LCA	1	0.00		
MAF	1	0.00		
SPM	1	0.00		
VCT	2	0.00		
WSM	4	0.00		
SMR	1	0.00		
STP	2	0.00		
SAU	538	24.35	256.89	2415.68
SEN	50	4.00	150.93	188.75
SRB	56	8.93	269.40	477.90
SYC	1	0.00		
SLE	11	0.00		
SGP	1	100.00	751.58	751.58
SXM	1	0.00		
SVK	40	0.00		
SVN	18	5.56	184.04	184.04
SLB	6	0.00		
SOM	39	0.00		
ZAF	446	7.85	181.15	725.13
SGS	5	0.00		
KOR	77	32.47	524.56	1062.34
ESP	286	30.42	285.45	948.28
S–	1	0.00		
LKA	32	6.25	177.83	205.81
SDN	264	2.27	141.46	341.84
SUR	21	9.52	7.60	8.90
SJM	0			
SWZ	10	0.00		
SWE	354	22.88	225.21	692.21
CHE	36	22.22	212.52	304.96
SYR	83	18.07	288.29	766.55
TWN	32	28.13	590.54	787.49
TJK	42	0.00		
TZA	154	0.65	141.88	141.88
THA	244	10.66	228.87	1038.65
TLS	11	0.00		
TGO	18	5.56	85.99	85.99
TKL	0			
TON	6	0.00		
TTO	6	83.33	299.55	417.19
ATF	0			
TUN	78	17.95	169.98	504.89
TUR	399	8.02	274.48	743.44
TKM	136	6.62	326.31	1288.06
TCA	4	0.00		
TUV	1	0.00		
UGA	42	4.76	130.44	132.34
UKR	360	3.89	157.89	323.22
ARE	43	48.84	385.80	1149.82
GBR	223	30.49	288.02	773.94
USA	3603	21.59	184.07	1904.87
URY	77	3.90	239.68	407.15
UZB	138	7.97	188.02	484.71
VAT	1	100.00	473.26	473.26
VEN	208	17.79	62.54	427.90
VNM	166	4.82	226.49	666.56
VIR	3	66.67	103.48	111.41
UMI	1	0.00		
WLF	3	0.00		
ESH	23	4.35	204.21	204.21
YEM	124	5.65	264.01	539.51
ZMB	105	3.81	106.12	204.82
ZWE	96	1.04	108.37	108.37
ZWE	0			

Table A-4: Light Gini with a Pareto Tail for the Top 1%, Germany, Pixel Level

Satellite	Gini unadjusted	α	y_c	μ_P	G_P	μ_R	G_R	G	Δ
F101992	0.3848	2.83	59	13.26	0.3703	91.21	0.2144	0.3979	0.0131
F101993	0.4038	2.83	58	11.43	0.3862	89.67	0.2144	0.4178	0.0140
F101994	0.3964	2.83	58	11.42	0.3784	89.67	0.2144	0.4107	0.0143
F121994	0.3607	2.83	60	15.56	0.3489	92.76	0.2144	0.3727	0.0120
F121995	0.4040	2.83	60	13.77	0.3913	92.76	0.2144	0.4165	0.0125
F121996	0.3913	2.83	60	13.93	0.3784	92.76	0.2144	0.4042	0.0129
F121997	0.3902	2.73	59	13.87	0.3772	93.08	0.2241	0.4033	0.0132
F121998	0.3800	2.63	60	15.85	0.3698	96.79	0.2346	0.3931	0.0131
F121999	0.3716	2.53	62	17.04	0.3622	102.50	0.2462	0.3855	0.0139
F141997	0.4173	2.73	59	11.70	0.4010	93.08	0.2241	0.4320	0.0147
F141998	0.4051	2.63	58	13.01	0.3917	93.56	0.2346	0.4194	0.0143
F141999	0.4088	2.53	59	12.20	0.3934	97.54	0.2462	0.4253	0.0165
F142000	0.4146	2.77	60	12.72	0.4004	93.89	0.2202	0.4285	0.0138
F142001	0.4219	2.84	60	12.36	0.4073	92.63	0.2138	0.4354	0.0135
F142002	0.4103	2.91	59	13.99	0.3987	89.93	0.2077	0.4218	0.0114
F142003	0.4314	2.98	61	12.88	0.4184	91.88	0.2020	0.4437	0.0123
F152000	0.3849	2.77	59	13.58	0.3715	92.32	0.2202	0.3986	0.0137
F152001	0.3974	2.84	61	14.05	0.3848	94.17	0.2138	0.4104	0.0130
F152002	0.3885	2.91	59	13.87	0.3754	89.93	0.2077	0.4004	0.0119
F152003	0.4218	2.98	54	10.09	0.4027	81.33	0.2020	0.4341	0.0123
F152004	0.4352	3.11	55	9.52	0.4149	81.04	0.1914	0.4476	0.0124
F152005	0.4293	3.15	57	11.26	0.4133	83.53	0.1888	0.4404	0.0112
F152006	0.4499	3.18	57	10.62	0.4336	83.09	0.1862	0.4612	0.0113
F152007	0.4361	3.30	57	10.29	0.4176	81.81	0.1788	0.4472	0.0111
F152008	0.3640	3.41	62	17.08	0.3543	87.73	0.1719	0.3729	0.0089
F162004	0.3944	3.11	57	11.42	0.3764	83.99	0.1914	0.4062	0.0118
F162005	0.4318	3.15	57	10.45	0.4134	83.53	0.1888	0.4436	0.0118
F162006	0.4072	3.18	57	11.89	0.3915	83.09	0.1862	0.4181	0.0109
F162007	0.4080	3.30	59	12.27	0.3926	84.68	0.1788	0.4186	0.0106
F162008	0.4204	3.41	60	12.58	0.4064	84.90	0.1719	0.4305	0.0102
F162009	0.3966	3.52	59	13.28	0.3827	82.40	0.1655	0.4056	0.0090
F182010	0.3285	3.63	62	22.95	0.3235	85.54	0.1595	0.3350	0.0065
F182011	0.4050	3.63	60	14.14	0.3931	82.78	0.1595	0.4134	0.0084
F182012	0.3778	3.63	62	16.97	0.3687	85.54	0.1595	0.3858	0.0081
F182013	0.3875	3.63	60	14.49	0.3753	82.78	0.1595	0.3959	0.0084

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

Note(s): Computed 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. Based on the DMSP-OLS v4 stable lights data after applying the GADM boundaries to isolate the island of Sicily. The data are adjusted such that the origins of the satellite images align, which may cause a slight displacement in some images. Each grid cell is 30 arc seconds by 30 arc seconds. We do not project the data onto an equal area grid since we want to induce minimal distortions. Instead, we weight each pixel by its land area for the applications.